arXiv:1405.2718v1 [q-fin.MF] 12 May 2014

ARBITRAGE PRICING OF MULTI-PERSON GAME CONTINGENT CLAIMS

Ivan Guo and Marek Rutkowski*
School of Mathematics and Statistics
University of Sydney
NSW 2006, Australia

6 April 2014

Abstract

We introduce a class of financial contracts involving several parties by extending the notion of a two-person game option to a contract in which an arbitrary number of parties is involved and each of them is allowed to make a wide array of decisions at any time, not restricted to simply 'exercising the option'. The collection of decisions by all parties then determines the contract's settlement date as well as the terminal payoff for each party. We provide sufficient conditions under which a multi-person game option has a unique arbitrage price, which is additive with respect to any partition of the contract.

Keywords: multi-person game claim, arbitrage price, Snell envelope, optimal equilibrium

Mathematics Subjects Classification (2010): 91A06, 91A15, 60G40

*The research of Ivan Guo and Marek Rutkowski was supported under Australian Research Council's Discovery Projects funding scheme (DP120100895).

1 Introduction

In financial markets, almost all traded derivatives effectively involve only two parties, the issuer and the holder. Furthermore, in most cases, the only meaningful decision lies with the holder in the form of exercising the contract. Some exceptions to this include the two-person game options introduced by Kifer [24] in which the issuer also has decisions to make (see also Dolinsky and Kifer [6], Dolinsky et al. [7], Kallsen and Kühn [20], Kühn et al. [26], Kyprianou [27] and a review paper by Kifer [25]). From the practical perspective, one may refer to convertible bonds as real-world financial derivatives covered by the general concept of the two-person game option (see, for instance, Ayache et al. [1], Andersen and Buffum [2], Bielecki et al. [3], or Kallsen and Kühn [21]). No wonder that the existing mathematical theory of pricing financial derivatives is mainly concerned with bilateral deals. By contrast, our goal is to open a new avenue by designing an extended framework in which one may define and evaluate contingent claims involving several parties.

In a general version of a multi-person contingent claim introduced in this work, an arbitrary number of parties is involved and each of them is allowed to make a wide array of decisions at any time, not restricted to simply 'exercising the option'. Decisions of all parties then determine the settlement time of the claim, as well as their respective terminal payoffs. Since one may draw equivalences between the mechanisms of multi-person financial contracts and general stochastic games examined in Guo and Rutkowski [11, 12, 13], we propose the name *multi-person game contingent claims* or *multi-person game options*. We work throughout within the discrete or continuous-time arbitrage-free market model with the intention of adding multi-person game contingent claims to this market. The generalised framework proposed here is not restricted to one particular issue, but has instead several purposes. First, it allows for the creation and valuation of new complex derivatives involving multiple parties. Second, it provides fresh insights to the valuation of existing exotic derivatives. Finally, the framework may serve as a starting point to a unified approach of modelling financial derivatives in the presence of external factors, such as: credit risk, funding costs, and margin account (see, for instance, [4, 31]). For instance, one may ask how the fact that external funders can make meaningful decisions in their own interest may affect the valuation of a contract between either two or several counterparties.

The paper is organized as follows. In Section 2, we provide the definition of a multi-person game contingent claim, which is associated with a multi-player stochastic game of perfect information. The claim involves several *holders* of parts of a contract (dubbed *tranches*), as well as its *issuer*. Since each party is able to observe the actions of others, this information is reflected in their decisions and hedging portfolios. As a result, admissible trading strategies are redefined as *$\mathcal{H}$-predictable trading strategies*, which also involve reactions to the observable actions of all parties. An *enter-and-hold arbitrage* is defined in a natural way as the opportunity to make a guaranteed profit by holding tranches along with an $\mathcal{H}$-predictable trading strategy until the claim is settled.

Section 3 focusses on the pricing of a particular *combined tranche,* that is, a fixed collection of tranches. To establish the no-arbitrage bounds for a combined tranche, it suffices to consider a holder of that tranche and to postulate that the issuer is holding all remaining tranches of a contract. The main results in this section are Theorems 3.1 and 3.3, for the discrete and continuous-time case, respectively. They show that, under some sensible conditions on payoffs, no enter-and-hold arbitrage is possible for a fixed combined tranche if and only if its price lies within a particular interval. The bounds of this interval are then interpreted in terms of super-hedging strategies.

In Section 4, we consider in addition the interaction between different combined tranches. In particular, we demonstrate that prices of single tranches of a multi-person game option must be additive in order to avoid a second type of arbitrage, dubbed the *immediate reselling arbitrage*. We show that the additivity of prices also eliminates any price disagreement between the parties. Theorem 4.1 furnishes the necessary and sufficient conditions for the non-existence of both types of arbitrage opportunities. Theorem 4.2 proves that if optimal equilibria exist for the associated multi-player game, then all individual tranches of a game contract have unique arbitrage prices, which match the values of the game. Furthermore, if the payoffs satisfy a sub-zero-sum condition,

then every combined tranche has a unique arbitrage price, which is also additive. For the reader's convenience, the pertinent definitions and results from the game theory are collected in the appendix.

1.1 Multi-Person Contract with Puttable Tranches

For concreteness, let us consider a continuous-time complete market model of a single stock, e.g., the Black-Scholes model. Our goal is to examine arbitrage pricing of contracts involving several parties. One of them, called the *issuer*, plays a special role – he initiates ('sells') the contract and takes care of its settlement according to the rules specified below. All other m parties are referred to as the *holders* of *tranches* of a contract. We give below an example of a simple contract between the issuer and m holders. In Section 4.5, we show how to value this contract using no-arbitrage principles.

1.1.1 Two-Period Setup

We consider dates $0 < T_1 < T_2 = T$ and we assume that at time 0 each holder buys a tranche, formally represented by a call or put option of either European or American style. Options may have different strikes, but they share the same expiry date T, and they cannot be exercised before the date T_1. At time T_1, each holder has the right to give back (that is, to *put*) the tranche to the issuer. If the ith holder decides to put the tranche he owns, then he receives from the issuer an $\mathcal{F}_{T_1}$-measurable amount X_1^i specified in the contract. We denote the set of such holders by $\mathcal{E}$. If all holders decide to put the tranche to the issuer at time T_1, then he simply pays all amounts $X_1^1, \ldots, X_1^m$ to holders. However, if at least one holder decides to keep his tranche, then the contract stipulates that the issuer computes his *total deviation* (the 'total loss') at time T_1, as given by the sum $\sum_{i\in\mathcal{E}}(X_1^i - P_1^i)$ where P_1^i is the price of the ith option at time T_1, and the total loss of the issuer is covered by the holders who decided to keep their tranches.

For simplicity, we assume that the *redistribution* is split evenly amongst the holders who decided to keep the tranche, so that each of them should pay the issuer the amount of $\frac{1}{k}\sum_{i\in\mathcal{E}}(X_1^i - P_1^i)$ where $k = m - |\mathcal{E}|$ is the number of such holders, although more general redistribution rules can also be implemented (see [11, 12, 13]). We note that it is possible to choose strikes $K_1, \ldots, K_m$, classes of options (calls or puts), and the amounts $X_1^1, \ldots, X_1^m$ to ensure that it will be not 'optimal' for all holders to put their tranches to the issuer at time T_1. Specifically, the amounts $X_1^1, \ldots, X_1^m$ may be chosen small enough so that there exists at least one holder for which it would be suboptimal to put his tranche, even when all other holders would prefer to do so. It is natural to conjecture that in this situation the issuer's 'total valuation' problem will reduce to the valuation of individual options, since holders' decisions at time T_1 will not affect him. This means that the sum of prices of 'tranches' at time 0 should be equal to the sum of prices of options. However, the price at time 0 of a tranche does not necessarily coincide with the price of the corresponding option, due to the redistribution of payoffs at time T_1.

1.1.2 Multi-Period Setup

Let us now consider the dates $0 < T_1 < T_2 < \cdots < T_n < T_{n+1} = T$. Each holder buys a tranche represented by an option expiring at T, which cannot be exercised before T_n. To describe the mechanism of the contract, let us fix some $l = 1, 2, \ldots, n$. At time T_l, each current holder has the right to put the tranche to the issuer. If the ith holder decides to do so, then he receives a predetermined $\mathcal{F}_{T_l}$-measurable amount X_l^i. We denote the set of such holders by $\mathcal{E}_l$. If all tranches are given back to the issuer at time T_l, then he pays all amounts $X_l^1, \ldots, X_l^m$ to holders. If at least one holder decides to keep his tranche, then the issuer computes his 'total loss' at time T_l, which is given by the sum $\sum_{i\in\mathcal{E}_l}(X_l^i - P_l^i)$ where P_l^i is the *market price* (*continuation value*) of the ith tranche at time T_l. The price P_l^i of tranche i is computed assuming that no putting of any tranche occurs at T_l (in other words, before holders decisions at time T_l were announced). The total loss of the issuer is covered by the holders who decided to keep their tranches. The redistribution is

split evenly amongst the holders who decided to keep the tranche at T_l, meaning that each of them compensate the issuer by paying him at time T_l the amount of $\frac{1}{k}\sum_{i\in\mathcal{E}_l}(X_l^i - P_l^i)$ where $k = m - |\mathcal{E}_l|$ is the number of such holders. After the redistribution stage has been completed, we move on to the next date T_{l+1}, with the issuer being the new holder of tranche i for all $i \in \mathcal{E}_l$. Then he can make any decisions available to the holder of tranche i from time T_{l+1} onwards. He can also sell some or all of these tranches to new holders and thus relinquish the corresponding decision-making abilities. We repeat these steps until the decisions and redistributions at date T_n are completed. At this moment, the tranches revert back to original options.

2 Multi-Person Game Options

All of the following definitions refer to an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which is endowed with the filtration $\mathbb{F} = \{\mathcal{F}_t : t \in [0,T]\}$ representing the information flow observed by all market participants. In continuous-time setup, the filtration $\mathbb{F}$ is also assumed to satisfy the 'usual conditions' that underpin stochastic calculus. We denote by $B, S^1, \ldots, S^d$ the $\mathbb{F}$-adapted processes representing prices of primary traded assets. For simplicity of presentation, it is postulated throughout that the underlying market model (which is left largely unspecified) is frictionless, arbitrage-free, and complete. Hence, according to the classic version of FTAP, the unique martingale measure, denoted by $\mathbb{Q}$, is available. For a detailed description of discrete-time case (continuous-time case, resp.), we refer to Jacod and Shiryaev [17] (Jeanblanc et al. [18], resp.). Further technical assumptions about the continuous time model will be stated in Section 3.5.

2.1 Actions, Outcomes and Strategy Profiles

Consider an arbitrary game $\mathcal{G}$ with player set given by $\mathcal{M} = \{1, 2, \ldots, m\}$ and the space of *actions* $\mathcal{A} = \prod_{i\in\mathcal{M}} \mathcal{A}^i$. We assumed that each $\mathcal{A}^i$ is a finite (or infinite, but countable) set representing actions available to player i at any time. Unless otherwise stated, the definitions cover both discrete- and continuous-time cases. In particular, $[t,u]$ denotes the set $\{t,\ldots,u\}$ when a discrete-time framework is studied. Note also that the convention $[0,t) = [0,t-1] := \{0,\ldots,t-1\}$ is used in that case for $t = 1,2,\ldots,T$.

Definition 2.1. An *outcome* for party i is an $\mathbb{F}$-adapted process $h^i = (h_t^i : t \in [0,T])$ with values in $\mathcal{A}^i$. It represents the actions of party i over time. The space of outcomes for party i is denoted by $\mathcal{H}^i$. We write $h = (h^1, \ldots, h^m)$ and $\mathcal{H} = \prod_{i\in\mathcal{M}} \mathcal{H}^i$ to denote the m-tuple of outcomes and its space, respectively.

We will frequently consider a sequence of actions up to, and possibly including, a certain time $t \le T$. For any outcome $h \in \mathcal{H}$, the *history before time* t is the restriction of h to the time interval $[0,t)$. We denote it by $h_{[0,t)}$ and we write $\mathcal{H}_{[0,t)}$ to denote the space of such histories. Also, we define $h_{[0,t]}$ and $\mathcal{H}_{[0,t]}$ in a similar way. The following technical assumption is relevant in the continuous-time case only, in which case it will be indeed crucial.

Assumption 2.1. Every outcome $h^i \in \mathcal{H}^i$ for every $i \in \mathcal{M}$ is a right-continuous process.

Many aspects of a game only depend on the past history only. To quantify this notion, let us consider a mapping $f : \mathcal{H} \times [0,T] \times \Omega \to C$ with values in some space C. For any $h, h' \in \mathcal{H}$, we define the random time

$$\rho(h,h') := \inf\left\{t \in [0,T] : h_t \neq h'_t\right\} \wedge T. \tag{2.1}$$

Then $\rho(h,h')$ is an $\mathbb{F}$-stopping time, since the outcomes h and h' are $\mathbb{F}$-adapted. Obviously, for every $t \in [0,T]$, the event $\{h_{[0,t)} = h'_{[0,t)}\}$ coincides with $\{t \le \rho(h,h')\}$.

Definition 2.2. A mapping $f : \mathcal{H} \times [0,T] \times \Omega \to C$ is *$\mathcal{H}$-adapted* (*$\mathcal{H}$-predictable*, resp.) if for every $h, h' \in \mathcal{H}$, the equality $f(h) = f(h')$ holds on the stochastic interval $[\![0,\rho)\!)$ (on $[\![0,\rho]\!]$, resp.) where $\rho = \rho(h,h')$.

It is clear that if a mapping f is $\mathcal{H}$-predictable then it is also $\mathcal{H}$-adapted. Intuitively, a mapping f is $\mathcal{H}$-adapted ($\mathcal{H}$-predictable, resp.) if, for all $t \in [0,T]$, the value $f_t(h)$ depends only on $h_{[0,t]}$ (on $h_{[0,t)}$, resp.). Informally, for an $\mathcal{H}$-adapted ($\mathcal{H}$-predictable, resp.) mapping f one could thus write $f_t(h) = f_t(h_{[0,t]})$ ($f_t(h) = f_t(h_{[0,t)})$, resp.) for every $h \in \mathcal{H}$ and $t \in [0,T]$.

A *strategy* for player i specifies how she should react to the observable actions of other parties. It is mapping from the space $\mathcal{H}$ of outcomes to the set $\mathcal{A}^i$ of actions, which also takes into account the information flow represented by the filtration $\mathbb{F}$.

Definition 2.3. A *strategy* of party i is an $\mathbb{F}$-adapted and $\mathcal{H}$-predictable mapping $s^i : \mathcal{H} \times [0,T] \times \Omega \to \mathcal{A}^i$. The m-tuple of strategies $s = (s^1, \ldots, s^m)$ is called a *strategy profile*. We denote by $\mathcal{S}^i$ ($\mathcal{S}$, resp.) the space of all possible strategies of party i (strategy profiles, resp.).

We define the mappings $s_{[0,t)} \in \mathcal{S}_{[0,t)}$ and $s^i_{[0,t)} \in \mathcal{S}^i_{[0,t)}$ in an obvious way. Note also that wherever the superscript $i \in \mathcal{M}$ appears, it can be replaced by an arbitrary subset $\mathcal{N} \subseteq \mathcal{M}$ to describe the corresponding notion relating to $\mathcal{N}$. Moreover, $-\mathcal{N}$ is a shorthand notation for $\mathcal{M} \setminus \mathcal{N}$. It will be frequently convenient to write

$$s = (\tau, \sigma) = (s^{\mathcal{N}}, s^{-\mathcal{N}}) \in \mathcal{S}^{\mathcal{N}} \times \mathcal{S}^{-\mathcal{N}}. \tag{2.2}$$

Definition 2.4. For any strategy profile $s \in \mathcal{S}$, we denote by $h(s) \in \mathcal{H}$ the *outcome associated* to s, that is, the unique outcome obtained when a strategy profile s is played on $[0,T]$.

It is natural to define a game by first specifying the class $\mathcal{S}^i$ of all strategies available to each player and, subsequently, to define $\mathcal{H}^i$ and $\mathcal{H}$ as the sets of all outcomes associated with all possible strategy profiles $s \in \mathcal{S}$. This specification underpins all definitions and results stated in what follows.

Let us now consider the game at time t. We assume that at any date t the sample path of the history $h_{[0,t)}(s)$ (but not the mapping $s_{[0,t)}$ used to generate this sample path) is observed by all parties. For any fixed $s = (\tau, \sigma) \in \mathcal{S}$ and $\mathcal{N} \subseteq \mathcal{M}$, we define the set $\mathcal{S}^{\mathcal{N}}_t(s)$ of strategies available to players from the set $\mathcal{N}$ at time t after a strategy profile $s_{[0,t)} \in \mathcal{S}_{[0,t)}$ was played. Formally, the subset $\mathcal{S}^{\mathcal{N}}_t(\tau, \sigma) \subseteq \mathcal{S}^{\mathcal{N}}$ is given by

$$\mathcal{S}^{\mathcal{N}}_t(\tau, \sigma) := \big\{\tau' \in \mathcal{S}^{\mathcal{N}} : h_{[0,t)}(\tau', \sigma) = h_{[0,t)}(\tau, \sigma)\big\}. \tag{2.3}$$

2.2 Multi-Person Game Contingent Claims

We first formulate a general definition of a multi-person game contract. An important ingredient of a contract is a decision to terminate it, which may be chosen by any party i provided, of course, that this possibility belongs to the action set $\mathcal{A}^i$. In an abstract formulation of a game claim, it will be enough to introduce the concept of the *termination mapping* through the following definition.

Assumption 2.2. The *termination mapping* $\Theta : \mathcal{H} \times [0,T] \times \Omega \to \{0,1\}$ satisfies:
(i) for every $h \in \mathcal{H}$, the process $\Theta(h)$ is $\mathbb{F}$-adapted, non-decreasing and RCLL,
(ii) the mapping Θ is $\mathcal{H}$-adapted.

For any $h \in \mathcal{H}$, the termination date is defined by

$$\vartheta(h) = \inf\{t \in [0,T] : \Theta_t(h) = 1\}. \tag{2.4}$$

Note that $\vartheta(h) \in \mathcal{T}_{[0,T]}$ for every $h \in \mathcal{H}$, where $\mathcal{T}_{[0,T]}$ stand for the class of all $\mathbb{F}$-stopping times with values in $[0,T]$. Condition (ii) in Assumption 2.2 means that for every $h, h' \in \mathcal{H}$, if $h_{[0,t]} = h'_{[0,t]}$ and $\vartheta(h) = t$ then $\vartheta(h') = t$. Given a strategy profile s and the associated outcome $h(s)$, we denote by $\vartheta(s)$ the termination date $\vartheta(h(s))$. Let $V_{\vartheta(s)}(s) = (V^1_{\vartheta(s)}(s), \ldots, V^m_{\vartheta(s)}(s))$ stand for the vector of random payoffs, which are received by holders at time $\vartheta(s)$ if a strategy profile s is played or, equivalently, if the outcome $h(s)$ is realised. Each particular payoff $V^i_{\vartheta(s)}(s)$ will be identified with a single *tranche* of a contract.

Assumption 2.3. The *payoff mapping* $V : \mathcal{H} \times [0,T] \times \Omega \to \mathbb{R}^m$ satisfies:
(i) for every $h \in \mathcal{H}$, the processes $V^1(h), \ldots, V^m(h)$ are $\mathbb{F}$-adapted and RCLL,
(ii) the mapping V is $\mathcal{H}$-adapted.

In the classic case of a two-person game option, as defined in Kifer [24], one deals with two tranches, which are held by a *holder* and an *issuer*, respectively. Note that, despite being attributed different names, their roles and rights are here identical. It should thus be stressed that in a multi-person game claim the role of the *issuer* will be essentially different from that of ordinary holders.

We use the symbol $\mathcal{G}$ to denote a generic game that underpins the specification of a multi-person game contingent claim. For explicitly defined classes of multi-player stochastic stopping games and results on the existence of an *optimal equilibrium*, we refer to [11, 12, 13]. The definition and basic properties of an optimal equilibrium are also given in the appendix.

Definition 2.5. A *multi-person game contingent claim* $\mathrm{MGC}(\mathcal{G})$ is a contract on $[0,T]$, which consists of m *tranches*, between an *issuer* and up to m *holders*. At time 0, a prospective holder of tranche i has to pay the issuer some fee (possibly negative) to enter the contract; we denote this amount by $\pi_0^i(\mathcal{G})$. Any individual party, including the issuer, is allowed to hold any combination of tranches. Once the contract starts, the holder of tranche i may choose a strategy $s^i \in \mathcal{S}^i$, and decisions of all holders form a strategy profile $s = (s^1, \ldots, s^m)$. The contract is settled at time $\vartheta(s) \in \mathcal{T}_{[0,T]}$ where the issuer pays the holder of tranche i the amount of $V^i_{\vartheta(s)}(s)$ for all $i = 1, \ldots, m$.

In the real-world financial contracts, all cashflows are zero-sum in nature. Since a game $\mathcal{G}$ is not postulated to be zero-sum, the crucial role of the issuer in $\mathrm{MGC}(\mathcal{G})$ is to absorb any surpluses or cover any deficits at times 0 and $\vartheta(s)$. Consequently, the net cashflow between the holders of the m tranches and the issuer is always zero. If $\mathcal{G}$ happens to be zero-sum, meaning that $\sum_{i\in\mathcal{M}} V_{\vartheta(s)}(s) = 0$ for all $s \in \mathcal{S}$, then the issuer is somewhat redundant, since the holders can strike the deal between themselves at time 0 and then also settle the payoffs between themselves at time $\vartheta(s)$ without issuer's intervention. In practice, the role of an issuers is then reduced to a clearing house or intermediation.

Remark 2.1. One may also reinterpret the issuer as the holder of an *issuer's tranche*, so that $\mathrm{MGC}(\mathcal{G})$ becomes a zero-sum game contract with $m+1$ tranches in which the holder of the $(m+1)$th tranche is not allowed to take any actions. This corresponds to the concept of a *dummy player* in [13]. Formally, the analysis of a game would then reduce to the zero-sum case.

Remark 2.2. The contract introduced in Definition 2.5 extends, in particular, the classic game option introduced in Kifer [24], which corresponds to a Dynkin stopping game between two players. The 'holder' owns the tranche, which allows for 'exercising', whereas the second holder (dubbed the 'issuer') keeps the tranche allowing for 'cancellation'. In the classic game options, there is no need to introduce an issuer in the sense of Definition 2.5, since the classic game option is associated with a zero-sum stopping game. Aside from the number and roles of parties involved, we also generalise the two-person game option by allowing a greater complexity in the underlying game $\mathcal{G}$. In particular, the analysis presented in what follows is by no means restricted to the case of stopping games.

2.3 Adaptive Trading Strategies

Let $\mathrm{MGC}_t(\mathcal{G}, s)$ denote the multi-person game contingent claim on $[t,T]$ after the strategy profile $s_{[0,t)} \in \mathcal{S}_{[0,t)}$ was played during $[0,t)$. Strictly speaking, $\mathrm{MGC}_t(\mathcal{G}, s)$ only depends on the history $h_{[0,t)}(s)$, meaning that if the equality $h_{[0,t)}(s) = h_{[0,t)}(s')$ holds, where s' is any strategy profile, then the game claims $\mathrm{MGC}_t(\mathcal{G}, s)$ and $\mathrm{MGC}_t(\mathcal{G}, s')$ are in fact identical. To simplify presentation, we implicitly work under the event $\{\vartheta(s) \geq t\}$ when dealing with $\mathrm{MGC}_t(\mathcal{G}, s)$. In other words, we only consider the case where $\mathrm{MGC}_t(\mathcal{G}, s)$ has not yet been terminated before time t.

Denote each *single tranche* $i \in \mathcal{M}$ by $\mathrm{MGC}^i(\mathcal{G})$ and the collection of single tranches $\mathcal{N} \subseteq \mathcal{M}$, dubbed a *combined tranche* $\mathcal{N}$, by $\mathrm{MGC}^{\mathcal{N}}(\mathcal{G})$. We will simply use the term *tranche* when no ambiguity may arise. For each $i \in \mathcal{M}$ and $\mathcal{N} \subseteq \mathcal{M}$, let $\mathrm{MGC}^i_t(\mathcal{G}, s)$ and $\mathrm{MGC}^{\mathcal{N}}_t(\mathcal{G}, s)$ denote the tranche i and the tranche $\mathcal{N}$ at time t, respectively.

Let us now describe trading arrangements after the contract's initiation. At any time t before the settlement time, all tranches of $\mathrm{MGC}_t(\mathcal{G}, s)$ may be freely traded, so that the holder of the tranche i may sell it to a new holder. The new holder then continues the game with the other current holders and has the ability to choose a strategy $\tau \in \mathcal{S}_t^i(s)$ from time t onwards. Denote the time t price of the tranche i (tranche $\mathcal{N}$, resp.) by $\pi_t^i(\mathcal{G}, s)$ ($\pi_t^{\mathcal{N}}(\mathcal{G}, s)$, resp.). Let us stress that anyone, including the issuer, is allowed to hold any number of tranches at any time.

The pricing of game options hinges on no-arbitrage arguments and the replication (or super-hedging) of payoffs. In the current framework, every party (the issuer and holders) may hold a portfolio of stock and bonds in order to hedge her future liabilities. Since $\mathcal{G}$ is a perfect information game, in which actions exist before $\vartheta(s)$ and are observed by all parties, it is natural for the issuer and the holders to adapt their trading strategies according to their current and past observations. In particular, if a strategy profile $s \in \mathcal{S}$ is played, a trading strategy should have a non-anticipative dependency on the observable outcome $h(s) \in \mathcal{H}$. This leads to the following definition in which it is implicitly assumed that the class of *admissible* (thus, in particular, self-financing) trading strategies for a given market model (B, S) was already specified (for instance, by focussing on trading strategies for which the discounted wealth is a $(\mathbb{Q}, \mathbb{F})$-martingale).

Definition 2.6. An *$\mathcal{H}$-predictable trading strategy* is a mapping $\phi : \mathcal{H} \times [0,T] \times \Omega \to \mathbb{R}^{d+1}$ satisfying the following conditions:
(i) for every $h \in \mathcal{H}$, $\phi(h)$ is an admissible trading strategy,
(ii) ϕ is an $\mathcal{H}$-predictable mapping.

The class of all $\mathcal{H}$-predictable trading strategies ϕ is denoted by $\Phi(\mathcal{H})$. For any $\phi \in \Phi(\mathcal{H})$, the corresponding wealth mapping $Z : \mathcal{H} \times [0,T] \times \Omega \to \mathbb{R}$, which is given by the equality

$$Z_t(\phi(h)) = \sum_{l=0}^{d} \phi_t^l(h) S_t^l, \quad \forall\, t \in [0,T],$$

is also $\mathcal{H}$-predictable. We will need to extend Definition 2.6 a little further. For a fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$, we denote by $\mathcal{H}^\sigma$ the set of all outcomes reachable when a strategy σ is played, that is,

$$\mathcal{H}^\sigma := \big\{ h \in \mathcal{H} : h = h(\tau, \sigma),\ \tau \in \mathcal{S}^{\mathcal{N}} \big\}. \tag{2.5}$$

In some situations, where possible outcomes are assumed to belong to $\mathcal{H}^\sigma$, we may be only interested in trading strategies with the same constraint and restricted to the interval $[t,T]$.

Definition 2.7. For any $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $t \in [0,T]$, we denote by $\Phi_t(\mathcal{H}^\sigma)$ the class of all *$\mathcal{H}^\sigma$-predictable trading strategies* $\phi : \mathcal{H}^\sigma \times [t,T] \times \Omega \to \mathbb{R}^{d+1}$. The classes $\mathcal{H}^\tau$ and $\Phi_t(\mathcal{H}^\tau)$ are defined in an analogous manner.

Intuitively, Definition 2.7 has the following interpretation: when a strategy profile σ is fixed then, for every $u \in [t,T]$ and $\tau, \tau' \in \mathcal{S}^{\mathcal{N}}$, the equality $h_{[0,u)}(\tau, \sigma) = h_{[0,u)}(\tau', \sigma)$ implies that $Z_u(\phi(\tau, \sigma)) = Z_u(\phi(\tau', \sigma))$. One could thus informally write $Z_u(\phi(\tau, \sigma)) = Z_u(\phi(\tau_{[0,u)}, \sigma))$ for every $\tau \in \mathcal{S}^{\mathcal{N}}$ and $u \in [t,T]$. In the continuous time framework, the latter property will underpin the definition of an $(\widehat{S}, \mathcal{H})$-predictable mapping ϕ (see Section 3.5).

3 Arbitrage Bounds for a Combined Tranche

An important concept in multi-person game contingent claims is the ability to hold a collection of tranches $\mathcal{N} \subseteq \mathcal{M}$. For example, a person holding tranches i and j simultaneously has possibly greater payoff earning potential than two separate people holding tranches i and j without collusion. This is due to the fact that the holder of tranches i and j may coordinate the strategies s^i and s^j to improve the combined payoff $V^i_{\vartheta(s)}(s) + V^j_{\vartheta(s)}(s)$. Consequently, the price of the combined tranche $\mathcal{N} = \{i, j\}$

is not necessarily equal to the sum of individual prices, even though from the practical perspective this additivity property seems to be desirable. The issue of price additivity and consistent valuation of all tranches will be further explored in Section 4.

In this section, we fix $\mathcal{N}$ and we restrict our attention to the pricing of a predetermined combined tranche $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$ for some fixed $t \in [0, T]$ and a strategy profile $s \in \mathcal{S}$. We address the following problem, in which we implicitly work on the event of $\{\vartheta(s) \geq t\}$ (since otherwise the question would be meaningless anyway):

- Consider the game contingent claim $\mathrm{MGC}(\mathcal{G})$ and suppose that a strategy profile $s_{[0,t)}$ has been played up to time t. At time t, a new holder purchases the tranche $\mathcal{N}$ from its previous holder(s) for the price of $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ and plans to hold it until the settlement time. What possible values can $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ take if neither issuer's nor holder's arbitrage in $(B, S, \mathcal{N})$ may occur?

3.1 Arbitrage Opportunities

We first define arbitrage opportunities from the perspective of a holder of a fixed combined tranche $\mathcal{N}$, as well as the issuer who is assumed to hold all remaining single tranches, that is, a combined tranche $-\mathcal{N}$. Recall the notation $s^{\mathcal{N}} = (s^i,\, i \in \mathcal{N})$ and $V_{\vartheta(s)}^{\mathcal{N}}(s) = \sum_{i \in \mathcal{N}} V_{\vartheta(s)}^i(s)$ for any subset $\mathcal{N} \subseteq \mathcal{M}$. We stress that in the rest of this section, a date t, a combined tranche $\mathcal{N}$ and a strategy profile s generating the history via $h_{[0,t)}(s)$ are fixed. Our goal in this section is to derive arbitrage bounds at time t for a fixed combined tranche $\mathcal{N}$. Nevertheless, the specification of payoffs for tranches from the set $-\mathcal{N}$ will still matter, since decisions of the holder of $-\mathcal{N}$ may impact the payoffs of tranches from $\mathcal{N}$.

In the next definition, we fix $t \in [0, T]$ and we assume that the tranche $\mathcal{N}$ of the game contract $\mathrm{MGC}_t(\mathcal{G}, s)$ is traded at time t at the price of $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$. To alleviate notation, we will simply write $(B, S, \mathcal{N})$ instead of a more precise notation $(B, S, \pi_t^{\mathcal{N}}(\mathcal{G}, s))$ if no confusion may arise.

Definition 3.1. A *holder's arbitrage in* $(B, S, \mathcal{N})$ is a pair $(\phi^\tau, \tau) \in \Phi_t(\mathcal{H}^\tau) \times \mathcal{S}_t^{\mathcal{N}}(s)$ for which there exists an event $A \in \mathcal{F}_t$ with a positive probability such that on A:

$$Z_t(\phi^\tau(s)) < -\pi_t^{\mathcal{N}}(\mathcal{G}, s), \quad Z_{\vartheta(s)}(\phi^\tau(s)) \geq -V_{\vartheta(s)}^{\mathcal{N}}(\tau, \sigma), \quad \forall\, \sigma \in \mathcal{S}_t^{-\mathcal{N}}(s).$$

An *issuer's arbitrage in* $(B, S, \mathcal{N})$ is a pair $(\phi^\sigma, \sigma) \in \Phi_t(\mathcal{H}^\sigma) \times \mathcal{S}_t^{-\mathcal{N}}(s)$ for which there exists an event $A \in \mathcal{F}_t$ with a positive probability such that on A:

$$Z_t(\phi^\sigma(s)) < \pi_t^{\mathcal{N}}(\mathcal{G}, s), \quad Z_{\vartheta(s)}(\phi^\sigma(s)) \geq V_{\vartheta(s)}^{\mathcal{N}}(\tau, \sigma), \quad \forall\, \tau \in \mathcal{S}_t^{\mathcal{N}}(s).$$

We say that there is *no arbitrage in* $(B, S, \mathcal{N})$ if neither the issuer's nor the holder's arbitrage exists in $(B, S, \mathcal{N})$.

Note that Definition 3.1 hinges on an implicit assumption that tranches are traded at time t before actions at time t are chosen by all parties. In other words, only the history $h_{[0,t)}$ and the σ-field $\mathcal{F}_t$ are observed by all parties before trading occurs at time t.

A holder's arbitrage refers to a guaranteed profit for the buyer of the tranche $\mathcal{N}$ on $[\![t, \vartheta(s)]\!]$. The price at time t should be high enough, so that a prospective holder should not be able to generate profits without risk by combining the buy-and-hold strategy with some clever actions and a dynamic portfolio composed of traded primary assets. Similarly, an issuer's arbitrage refers to a guaranteed profit for issuer of $\mathrm{MGC}(\mathcal{G})$ who sold the tranche $\mathcal{N}$ at time t and also holds the tranche $-\mathcal{N}$ until time $\vartheta(s)$. From this perspective, the price should be low enough, so that it prevents the issuer from making profit without risk even when he can make decisions regarding the tranche $-\mathcal{N}$. When the issuer does not hold the tranche $-\mathcal{N}$, we may interpret an issuer's arbitrage as a guaranteed loss for the holder of the tranche $\mathcal{N}$. It is rather clear that, under Definition 3.1, the valuation problem for the tranche $\mathcal{N}$ can be formally reduced to a study of the two-person zero-sum game between the holder of the tranche $\mathcal{N}$ and the issuer, who also holds $-\mathcal{N}$.

Arbitrages introduced in Definition 3.1 hinge on the enter-and-hold strategy, where the tranches are kept from time t until the settlement time. Of course, alternative definitions of arbitrage opportunities are also possible. In Section 4, we will complement Definition 3.1 with additional conditions that will ensure the existence and uniqueness of financially meaningful unique prices for all tranches.

3.2 Super-Hedging Strategies

For a European contingent claim, the terminal payoff can be frequently *replicated* by means of an admissible portfolio. In the case of a game contingent claim, since it is not possible to anticipate the action of other parties, replication is typically not possible. Instead, the parties involved may attempt to *super-hedge* their positions. The definition of super-hedging trading strategies is very similar to that for the two-person game option. As before, we consider the time t game contingent claim $\mathrm{MGC}_t(\mathcal{G}, s)$. We focus on the holder of the tranche $\mathcal{N}$ and the issuer, who also holds the tranche $-\mathcal{N}$. Recall that we denote $s = (\tau, \sigma)$ and we place ourselves at time $t \in [0, T]$.

Definition 3.2. A *holder's super-hedging strategy* is a pair $(\phi^\tau, \tau) \in \Phi_t(\mathcal{H}^\tau) \times \mathcal{S}_t^{\mathcal{N}}(s)$ such that

$$Z_{\vartheta(s)}(\phi^\tau) \geq -V^{\mathcal{N}}_{\vartheta(s)}(\tau, \sigma), \quad \forall\, \sigma \in \mathcal{S}_t^{-\mathcal{N}}(s).$$

An *issuer's super-hedging strategy* is a pair $(\phi^\sigma, \sigma) \in \Phi_t(\mathcal{H}^\sigma) \times \mathcal{S}_t^{-\mathcal{N}}(s)$ such that

$$Z_{\vartheta(s)}(\phi^\sigma) \geq V^{\mathcal{N}}_{\vartheta(s)}(\tau, \sigma), \quad \forall\, \tau \in \mathcal{S}_t^{\mathcal{N}}(s).$$

As expected, the super-hedging strategies are closely related to arbitrage bounds for prices of tranches of game contingent claims.

Proposition 3.1. *Suppose that the tranche $\mathcal{N}$ of the game contract $\mathrm{MGC}_t(\mathcal{G}, s)$ is traded at time t at a price of $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$. If no issuer's arbitrage in $(B, S, \mathcal{N})$ exists, then for any issuer's super-hedging strategy $(\phi^\sigma, \sigma) \in \Phi_t(\mathcal{H}^\sigma) \times \mathcal{S}_t^{-\mathcal{N}}(s)$ we have that*

$$Z_t(\phi^\sigma) \geq \pi_t^{\mathcal{N}}(\mathcal{G}, s). \tag{3.1}$$

If no holder's arbitrage in $(B, S, \mathcal{N})$ exists, then for any holder's super-hedging strategy $(\phi^\tau, \tau) \in \Phi_t(\mathcal{H}^\tau) \times \mathcal{S}_t^{\mathcal{N}}(s)$ we have that

$$Z_t(\phi^\tau) \geq -\pi_t^{\mathcal{N}}(\mathcal{G}, s). \tag{3.2}$$

Proof. Let us show that if the upper bound (3.1) is violated, then there exists an issuer's arbitrage. A similar argument can be used to check that the violation of the lower bound (3.2) leads to a holder's arbitrage. For the sake of contradiction, let us assume the upper bound of (3.1) is not satisfied. On the event $\{\pi_t^{\mathcal{N}}(\mathcal{G}, s) > Z_t(\phi^\sigma)\}$, by the definition of an issuer's super-hedging strategy, we obtain

$$Z_t(\phi^\sigma) < \pi_t^{\mathcal{N}}(\mathcal{G}, s), \quad Z_{\vartheta(s)}(\phi^\sigma) \geq V^{\mathcal{N}}_{\vartheta(s)}(\tau, \sigma), \quad \forall\, \tau \in \mathcal{S}_t^{\mathcal{N}}(s).$$

Hence (ϕ^σ, σ) is an issuer's arbitrage in $(B, S, \mathcal{N})$, which contradicts our assumption. □

3.3 $\mathcal{H}^\sigma$-Predictable Snell Envelopes

The notion of the Snell envelope, first introduced in [33], is an important tool for the valuation of American and game options. For a given stochastic process Y, the *Snell envelope* U is the smallest RCLL supermartingale dominating Y. At time t, the Snell envelope U_t equals $\operatorname{ess\,sup}_{\tau \in \mathcal{T}_{[t,T]}} \mathbb{E}(Y_\tau \,|\, \mathcal{F}_t)$, so that it represents the essential supremum of $\mathcal{F}_t$-conditional expectations of Y_τ over all choices of a stopping time $\tau \in \mathcal{T}_{[t,T]}$.

A similar notion will be developed here in the context of a multi-person game claim $\mathrm{MGC}(\mathcal{G})$. It will hinge on maximizing an expected value over the set of all possible strategies, rather than the

set of all stopping times. Recall that the subset $\mathcal{N} \in \mathcal{M}$ is now fixed and we deal with a combined tranche, specifically, a collection of single tranches i for all $i \in \mathcal{N}$. Unless otherwise stated, we fix a strategy profile $s = (\tau, \sigma) \in \mathcal{S}$ where $\tau \in \mathcal{S}^{\mathcal{N}}$ and $\sigma \in \mathcal{S}^{-\mathcal{N}}$.

Recall that we work throughout under the unique martingale measure $\mathbb{Q}$ for a discrete-time (or a continuous-time) market model, which is assumed to be complete and arbitrage-free in the usual sense. Let $\widehat{V}_{\vartheta(s)}(s) = B^{-1}_{\vartheta(s)} V_{\vartheta(s)}(s)$ be the discounted payoff and let $\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(s) = \sum_{i\in\mathcal{N}} \widehat{V}^{i}_{\vartheta(s)}(s)$. In the remainder of this work, we work under the following standing assumption.

Assumption 3.1. The discounted payoff satisfies the following integrability condition with respect to $\mathbb{Q}$, for every $\mathcal{N} \subseteq \mathcal{M}$,

$$\mathbb{E}_{\mathbb{Q}}\Big(\operatorname*{ess\,sup}_{s\in\mathcal{S}} \big|\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(s)\big|\Big) < \infty.$$

3.3.1 $\mathcal{H}^\sigma$-Predictable Snell Envelope of the First Kind

We first define the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind, which is directly based on the observations of decisions of the holder of the tranche $\mathcal{N}$ strictly before time t and the market data up to time t. Note that the holder of $-\mathcal{N}$ has a priori the full knowledge of her fixed strategy σ on $[0,T]$, but she only observes the outcome $h_{[0,t)}(\tau,\sigma)$ at time t.

Definition 3.3. For a combined tranche $\mathrm{MGC}^{\mathcal{N}}(\mathcal{G})$ and a fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$, the *$\mathcal{H}^\sigma$-predictable Snell envelope of the first kind* is the mapping $U^\sigma : \mathcal{H}^\sigma \times [0,T] \times \Omega \to \mathbb{R}$ given by

$$U^\sigma_t(\tau) = U^\sigma_t(h(\tau,\sigma)) := \operatorname*{ess\,sup}_{\tau'\in\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\,\big|\,\mathcal{F}_t\Big). \tag{3.3}$$

Remark 3.1. Let us note that the definition of $U^\sigma_t(\tau)$ is not restricted to the event $\{\vartheta(\tau,\sigma) \geq t\}$. Indeed, on the event $\{\vartheta(\tau,\sigma) < t\}$, by convention, we may set

$$U^\sigma_t(\tau) := \widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma) \leq U^\sigma_{\vartheta(s)}(\tau).$$

Recall that the class $\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)$ is defined by

$$\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma) = \big\{\tau' \in \mathcal{S}^{\mathcal{N}} : h_{[0,t)}(\tau',\sigma) = h_{[0,t)}(\tau,\sigma)\big\}. \tag{3.4}$$

Since σ and τ are here fixed, the class $\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)$ only depends on $h_{[0,t)}(\tau,\sigma)$. Therefore, the right-hand side in (3.3) is indeed a well-defined mapping of the outcome $h(\tau,\sigma) \in \mathcal{H}^\sigma$, rather than the strategy (τ,σ). Despite this, for the sake of brevity, we prefer to write $U^\sigma_t(\tau)$, rather than $U^\sigma_t(h(\tau,\sigma))$.

The superscript σ in $U^\sigma(\tau)$ is used to emphasise that the strategies σ and τ play different roles in Definition 3.3. Specifically, at any date t we assume to know an issuer's strategy σ over $[0,T]$, but we only assume that a holder's strategy τ is observed on $[0,t)$. This feature is formalised in the next result, which shows that, intuitively, $U^\sigma_t(\tau) = U^\sigma_t(\tau_{[0,t)})$ for all $t \in [0,T]$, where $\tau_{[0,t)}$ is the restriction of the holder's strategy τ to the interval $[0,t)$.

Lemma 3.1. *The mapping U^σ is $\mathcal{H}^\sigma$-predictable.*

Proof. It is obvious from (3.3) that $\tau \in \mathcal{S}^{\mathcal{N}}$ the random variable $U^\sigma_t(\tau)$ is $\mathcal{F}_t$-measurable. For the $\mathcal{H}^\sigma$-predictability of U^σ, we need to show that, for all $\tau,\tau' \in \mathcal{S}^{\mathcal{N}}$, the equality $U^\sigma(\tau) = U^\sigma(\tau')$ holds on $[\![0,\rho]\!]$ where (see (2.1))

$$\rho := \inf\big\{t \in [0,T] : h_t(\tau,\sigma) \neq h_t(\tau',\sigma)\big\} \wedge T. \tag{3.5}$$

Recall that, by Assumption 2.3, the payoff mapping $V^{\mathcal{N}}$ (and thus also the discounted payoff $\widehat{V}^{\mathcal{N}}$) is $\mathcal{H}$-adapted (though not necessarily $\mathcal{H}^\sigma$-predictable). For any $t \in [0,T]$, the event $\{t \leq \rho\}$ belongs to $\mathcal{F}_t$ and the sets

$$\big\{\widehat{V}^{\mathcal{N}}_{\vartheta(\widetilde{\tau},\sigma)}(\widetilde{\tau},\sigma)\mathbb{1}_{\{t\leq\rho\}} : \widetilde{\tau} \in \mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)\big\} \quad \text{and} \quad \big\{\widehat{V}^{\mathcal{N}}_{\vartheta(\widehat{\tau},\sigma)}(\widehat{\tau},\sigma)\mathbb{1}_{\{t\leq\rho\}} : \widehat{\tau} \in \mathcal{S}^{\mathcal{N}}_t(\tau',\sigma)\big\}$$

are identical, due to the definition of $\mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)$ and $\mathcal{S}_t^{\mathcal{N}}(\tau',\sigma)$ (see (3.4)). Consequently, the equality $U_t^\sigma(\tau) = U_t^\sigma(\tau')$ holds on $\{t \leq \rho\}$ for every t. □

The $\mathcal{H}^\sigma$-predictable Snell envelope U^σ has the following interpretation: suppose that the holder of $-\mathcal{N}$ plays the strategy $\sigma \in \mathcal{S}^{-\mathcal{N}}$ on $[0,T]$ while the holder of $\mathcal{N}$ plays the strategy $\tau \in \mathcal{S}^{\mathcal{N}}$ on $[0,t)$. Then $U_t^\sigma(\tau)$ is the maximal expected payoff, as seen at time t, which can be achieved by the holder of $\mathcal{N}$ by varying his strategy on $[t,T]$. It is thus natural to conjecture that the process $U_t^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale. The proof of the main result in this subsection, in which this conjecture is validated (Proposition 3.2), requires several lemmas.

Lemma 3.2. *Fix $t \in [0,T]$, $\mathcal{N} \subset \mathcal{M}$ and $s \in \mathcal{S}$. Then the following set $\mathcal{D}_t$ has the lattice property*

$$\mathcal{D}_t := \Big\{ \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t \Big) : \tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma) \Big\}.$$

Proof. Let I be an arbitrary set and let $\mathcal{Y} := (Y_i)_{i\in I}$ be a family of random variables with $\mathbb{E}_{\mathbb{Q}}(\operatorname{ess\,sup}_{i\in I} |Y_i|) < \infty$. Then $\mathcal{Y}$ is said to have the *lattice property* if, for all $i,j \in I$, there exist $k,l \in I$ such that

$$Y_k \geq Y_i \vee Y_j, \quad Y_l \leq Y_i \wedge Y_j.$$

In our case, the set $\mathcal{D}_t$ may be represented as $\{Y_{\tau'} : \tau' \in \mathcal{S}_t^{\mathcal{N}}(s)\}$ where each $Y_{\tau'}$ is $\mathcal{F}_t$-measurable and, in view of Assumption 3.1,

$$\mathbb{E}_{\mathbb{Q}}\big(\operatorname*{ess\,sup}_{\tau'\in\mathcal{S}_t^{\mathcal{N}}(s)} |Y_{\tau'}| \big) < \infty.$$

We first show that it has the following property: for any event $A \in \mathcal{F}_t$ and arbitrary $\tau',\tau'' \in \mathcal{S}_t^{\mathcal{N}}(s)$

$$\tau' \mathbb{1}_A = \tau'' \mathbb{1}_A \implies Y_{\tau'} \mathbb{1}_A = Y_{\tau''} \mathbb{1}_A. \tag{3.6}$$

In view of the $\mathcal{H}$-adaptedness of the settlement time ϑ and the discounted payoff $\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma)$, the equality $\tau' \mathbb{1}_A = \tau'' \mathbb{1}_A$ implies that

$$\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\mathbb{1}_A = \widehat{V}^{\mathcal{N}}_{\vartheta(\tau'',\sigma)}(\tau'',\sigma)\mathbb{1}_A.$$

Since $\mathbb{1}_A$ is $\mathcal{F}_t$-measurable, we obtain $Y_{\tau'}\mathbb{1}_A = Y_{\tau''}\mathbb{1}_A$, thus establishing (3.6). Let us now check the lattice property of $\mathcal{D}_t$. For arbitrary $\tau',\tau'' \in \mathcal{S}_t^{\mathcal{N}}(s)$, we define the strategies τ_1,τ_2 by setting

$$\tau_1 = \mathbb{1}_{\{Y_{\tau'} \geq Y_{\tau''}\}}\tau' + \mathbb{1}_{\{Y_{\tau'} < Y_{\tau''}\}}\tau'',$$
$$\tau_2 = \mathbb{1}_{\{Y_{\tau'} \geq Y_{\tau''}\}}\tau'' + \mathbb{1}_{\{Y_{\tau'} < Y_{\tau''}\}}\tau'.$$

Since the event $\{Y_{\tau'} \geq Y_{\tau'}\}$ is $\mathcal{F}_t$-measurable, it is clear that τ_1,τ_2 are $\mathbb{F}$-adapted, $\mathcal{H}$-predictable mappings satisfying

$$h_{[0,t)}(\tau_1,\sigma) = h_{[0,t)}(\tau_2,\sigma) = h_{[0,t)}(\tau',\sigma) = h_{[0,t)}(\tau'',\sigma) = h_{[0,u)}(s), \quad \forall\, \sigma \in \mathcal{S}_t^{-\mathcal{N}}(s).$$

Hence τ_1 and τ_2 belong to $\mathcal{S}_t^{\mathcal{N}}(s)$. Using condition (3.6), we also obtain $Y_{\tau_1} = Y_{\tau'} \vee Y_{\tau''}$ and $Y_{\tau_2} = Y_{\tau'} \wedge Y_{\tau''}$. This establishes the lattice property of the set $\mathcal{D}_t$. □

The following result is a minor extension of Theorem A.32 in [9] (for the proof, see Lemma 7.18 in [10]).

Lemma 3.3. *Let $(Y_i)_{i\in I}$ be a family of $\mathcal{F}_u$-measurable random variables with the lattice property and such that $\mathbb{E}(\operatorname{ess\,sup}_{i\in I} |Y_i|) < \infty$. Then the following statements are valid:*
(i) there exist two sequences $(i_n)_{n\in\mathbb{N}}, (j_n)_{n\in\mathbb{N}}$ of indices such that the sequences $(Y_{i_n})_{n\in\mathbb{N}}$ and $(-Y_{j_n})_{n\in\mathbb{N}}$ are almost surely non-decreasing and

$$\operatorname*{ess\,sup}_{i\in I} Y_i = \sup_{n\in\mathbb{N}} Y_{i_n} = \lim_{n\to\infty} Y_{i_n}, \quad \operatorname*{ess\,inf}_{j\in I} Y_j = \inf_{n\in\mathbb{N}} Y_{j_n} = \lim_{n\to\infty} Y_{j_n},$$

(ii) for any $t \leq u$, we have

$$\mathbb{E}_{\mathbb{Q}}\big(\operatorname*{ess\,sup}_{i\in I} Y_i \,\big|\, \mathcal{F}_t \big) = \operatorname*{ess\,sup}_{i\in I} \mathbb{E}\big(Y_i \,|\, \mathcal{F}_t\big), \quad \mathbb{E}_{\mathbb{Q}}\big(\operatorname*{ess\,inf}_{j\in I} Y_j \,\big|\, \mathcal{F}_t \big) = \operatorname*{ess\,inf}_{j\in I} \mathbb{E}\big(Y_j \,|\, \mathcal{F}_t\big).$$

Lemma 3.4. *(i) The following equality holds, for any $t \leq u$,*

$$\mathbb{E}_{\mathbb{Q}}\big(U_u^{\sigma}(\tau) \,\big|\, \mathcal{F}_t\big) = \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big). \tag{3.7}$$

(ii) For any $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $U^{\sigma}(\tau)$ is uniformly integrable under $\mathbb{Q}$.

Proof. (i) Let us consider $t \leq u$. From Lemma 3.2, we deduce that, for fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $\tau \in \mathcal{S}^{\mathcal{N}}$, the set

$$\mathcal{D}_u = \Big\{\mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_u\Big) : \tau' \in \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma)\Big\} \tag{3.8}$$

has the lattice property. Using part (ii) in Lemma 3.3 and the tower property of conditioning, we thus obtain

$$\mathbb{E}_{\mathbb{Q}}\big(U_u^{\sigma}(\tau) \,\big|\, \mathcal{F}_t\big) = \mathbb{E}_{\mathbb{Q}}\Big(\operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_u\Big) \,\Big|\, \mathcal{F}_t\Big) = \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big)$$

so that (3.7) is valid.

(ii) In view of Assumption 3.1, the process M given by $M_t = \mathbb{E}_{\mathbb{Q}}(M_T \,|\, \mathcal{F}_t)$, where

$$M_T := \operatorname*{ess\,sup}_{s \in \mathcal{S}} \big|\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(s)\big|,$$

is a non-negative and uniformly integrable $(\mathbb{Q},\mathbb{F})$-martingale. Using the conditional Fatou lemma, we obtain, for any $t \in [0,T]$,

$$\begin{aligned}\big|U_t^{\sigma}(\tau)\big| &= \Big|\operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big)\Big| \leq \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\big|\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\big| \,\Big|\, \mathcal{F}_t\Big) \\ &\leq \mathbb{E}_{\mathbb{Q}}\Big(\operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)} \big|\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\big| \,\Big|\, \mathcal{F}_t\Big) \leq \mathbb{E}_{\mathbb{Q}}\Big(\operatorname*{ess\,sup}_{s \in \mathcal{S}} \big|\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(s)\big| \,\Big|\, \mathcal{F}_t\Big) = M_t.\end{aligned}$$

Hence the process $U^{\sigma}(\tau)$ is uniformly integrable under $\mathbb{Q}$. □

Proposition 3.2. *(i) For any $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $U^{\sigma}(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale.*
(ii) In the continuous-time case, if the process $U^{\sigma}(\tau)$ has an RCLL modification, then it is a $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D).

Proof. (i) It is clear that $U^{\sigma}(\tau)$ is an $\mathbb{F}$-adapted process. The supermartingale property of $U^{\sigma}(\tau)$ follows immediately from part (i) in Lemma 3.4 and the obvious inclusion $\mathcal{S}_u^{\mathcal{N}}(\tau,\sigma) \subseteq \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)$ for all $t \leq u$, since

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}}\big(U_u^{\sigma}(\tau) \,\big|\, \mathcal{F}_t\big) &= \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big) \\ &\leq \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big) = U_t^{\sigma}(\tau).\end{aligned}$$

(ii) Recall that a process X is said to be of *class (D)* if it is RCLL and the family $\{X_\rho : \rho \in \mathcal{T}_{[0,T]}\}$ of random variables is uniformly integrable. We assumed that $U^{\sigma}(\tau)$ has an RCLL modification. Let M be an RCLL version of the martingale M. Then $|U_\rho^{\sigma}(\tau)| \leq M_\rho$ for any stopping time $\rho \in \mathcal{T}_{[0,T]}$ where the family $\{M_\rho : \rho \in \mathcal{T}_{[0,T]}\}$ is uniformly integrable. Hence the family $\{U_\rho^{\sigma}(\tau) : \rho \in \mathcal{T}_{[0,T]}\}$ is uniformly integrable and thus $U^{\sigma}(\tau)$ is a supermartingale of class (D). □

3.3.2 $\mathcal{H}^\sigma$-Adapted Snell Envelope

Definition 3.4 of the $\mathcal{H}^\sigma$-adapted Snell envelope at time t hinges on the assumption that the decisions of all parties at time t are already known. It other words, it is now assumed that, at time t, the holder of $-\mathcal{N}$ observes the market data up to time t, as well as the outcome $h_{[0,t]}(\tau,\sigma)$. It is thus rather clear that the $\mathcal{H}^\sigma$-adapted Snell envelope is not directly suitable for super-hedging purposes. However, it will serve as a crucial tool in the introduction of the $\mathcal{H}^\sigma$-predictable Snell envelope of the second kind in the continuous-time setup (see Section 3.5.2).

Definition 3.4. For a combined tranche $\mathrm{MGC}^{\mathcal{N}}(\mathcal{G})$ and a fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$, we define the $\mathcal{H}^\sigma$*-adapted Snell envelope* $\widetilde{U}^\sigma : \mathcal{H}^\sigma \times [0,T] \times \Omega \to \mathbb{R}$ by setting

$$\widetilde{U}^\sigma(\tau) = \widetilde{U}^\sigma_t(h(\tau,\sigma)) := \operatorname*{ess\,sup}_{\tau' \in \widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big) \tag{3.9}$$

where the subset $\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma) \subseteq \mathcal{S}^{\mathcal{N}}$ is given by

$$\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma) := \big\{\tau' \in \mathcal{S}^{\mathcal{N}} : h_{[0,t]}(\tau',\sigma) = h_{[0,t]}(\tau,\sigma)\big\}. \tag{3.10}$$

Lemma 3.5. *(i) The mapping $\widetilde{U}^\sigma$ is $\mathcal{H}^\sigma$-adapted.*
(ii) Fix $t \in [0,T]$, $\mathcal{N} \subset \mathcal{M}$ and $s \in \mathcal{S}$. Then the following set $\widetilde{\mathcal{D}}_t$ has the lattice property

$$\widetilde{\mathcal{D}}_t := \Big\{\mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big) : \tau' \in \widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)\Big\}.$$

Proof. The proof of part (i) (part (ii), resp.) is analogous to the proof of Lemma 3.1 (Lemma 3.2, resp.) and thus it is omitted. □

Proposition 3.3. *(i) The following equality holds, for any $t \le u$,*

$$\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_u(\tau) \,\big|\, \mathcal{F}_t\big) = \operatorname*{ess\,sup}_{\tau' \in \widetilde{\mathcal{S}}^{\mathcal{N}}_u(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big).$$

(ii) The process $\widetilde{U}^\sigma(\tau)$ is uniformly integrable under $\mathbb{Q}$. In particular,

$$\mathbb{E}_{\mathbb{Q}}\Big(\operatorname*{ess\,sup}_{t\in[0,T]} \big|\widetilde{U}^\sigma_t(\tau)\big|\Big) < \infty.$$

(iii) For any $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $\widetilde{U}^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale.

Proof. In view of Lemma 3.5, all statements can be established in a similar way as analogous statements for the mapping $U^\sigma(\tau)$ (see the proofs of Lemma 3.4 and Proposition 3.2). □

3.4 Discrete-Time Case

We now assume a discrete-time arbitrage-free and complete market model (B,S) with a finite underlying probability space Ω. Therefore, for any European contingent claim with the payoff X at its maturity date T, there exists a self-financing trading strategy ϕ satisfying, for all $t \in [0,T]$,

$$Z_t(\phi) := \sum_{i=0}^{d} \phi^i_t S^i_t = \mathbb{E}_{\mathbb{Q}}(X_T \,|\, \mathcal{F}_t).$$

Note that the portfolio ϕ_t is established at time $t-1$, so that it is $\mathcal{F}_{t-1}$-measurable, meaning that the process ϕ is $\mathbb{F}$-predictable, and ϕ_t may depend on the outcome $h_{[0,t-1]}$. At time t, the portfolio is revised from ϕ_t to ϕ_{t+1} such that

$$\sum_{i=0}^{d} \phi^i_t S^i_t = \sum_{i=0}^{d} \phi^i_{t+1} S^i_t,$$

ϕ_{t+1} is $\mathcal{F}_t$-measurable and it may depend on $h_{[0,t]}$. Since the model is arbitrage-free, all trading strategies that replicate X have the same wealth process. The uniqueness of a replicating strategy is not guaranteed, however, since the complete market (B, S) may still have redundancies.

We return to the question posed in the beginning of the section, that is, for a given tranche $\mathcal{N}$ we search for the possible prices of $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ for some fixed $t \in [0, T]$ and $s \in \mathcal{S}$ that avoid arbitrage in $(B, S, \mathcal{N})$. We will first focus on the discrete-time case, before proceeding to the continuous-time one. Recall the convention $[0, t) = [0, t-1] := \{0, \dots, t-1\}$ for every $t = 1, 2, \dots, T$. We denote by $\widehat{Z}(\phi)$ the *discounted wealth* given as $\widehat{Z}(\phi) = B^{-1}Z(\phi)$. We will need the following lemma.

Lemma 3.6. *Suppose that a mapping $\widehat{M} : \mathcal{H} \times [0, T] \times \Omega \to \mathbb{R}$ satisfies the following properties:*
(i) for any $s \in \mathcal{S}$, the process $\widehat{M}(s) = \widehat{M}(h(s))$ is a $(\mathbb{Q}, \mathbb{F})$-martingale,
(ii) $\widehat{M}$ is $\mathcal{H}$-predictable.
Then there exists an $\mathcal{H}$-predictable mapping ϕ such that, for every $s \in \mathcal{S}$, the discounted wealth of a trading strategy $\phi(s)$ satisfies $\widehat{Z}(\phi(s)) = \widehat{M}(s)$.

Proof. For any $s \in \mathcal{S}$, we simply define $\phi(s)$ as a replicating (hence necessarily admissible) trading strategy for the claim $B_T\widehat{M}_T(s)$. Hence the existence of $\phi(s)$ follows immediately from the postulated completeness of the discrete-time market model (B, S) and $\mathbb{Q}$-integrability of $\widehat{M}_T(s)$. It is also not hard to show that an $\mathcal{H}$-predictable version of the mapping ϕ can be selected (for details, see Section 7.2.2 in Guo [10]). □

Remark 3.2. Note that the statement of the lemma will still hold if we replace $\mathcal{H}$ and $[0, T]$ by $\mathcal{H}^\sigma$ and $[t, T]$, respectively.

The next result shows that, for any fixed $(\tau, \sigma) \in \mathcal{S}$ and $t \in [0, T]$, there exists an issuer's super-hedging strategy (ϕ^σ, σ) on $[t, T]$ whose discounted wealth at time t coincides with the value $U_t^\sigma(\tau)$ of the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind. This property is a essential tool in establishing the upper bound in Theorem 3.1.

Proposition 3.4. *Take an arbitrary $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $t \in [0, T]$. If a strategy $\tau \in \mathcal{S}^{\mathcal{N}}$ is played by the holder on $[0, t-1]$, then there exists a trading strategy $\phi^\sigma(\tau) \in \Phi_t(\mathcal{H}^\sigma)$ such that $(\phi^\sigma(\tau), \sigma)$ is an issuer's super-hedging strategy on $[t, T]$ and its discounted time t wealth is equal to the value $U_t^\sigma(\tau)$ of the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind. More explicitly, the discounted wealth process $\widehat{Z}_u(\phi^\sigma(\tau)) := B_u^{-1}Z_u(\phi^\sigma(\tau))$, $u \in [t, T]$, satisfies, on the event $\{\vartheta(\tau, \sigma) \geq t\}$,*

$$\widehat{Z}_t(\phi^\sigma(\tau)) = U_t^\sigma(\tau) \tag{3.11}$$

and

$$\widehat{Z}_{\vartheta(\tau,\sigma)}(\phi^\sigma(\tau)) \geq U^\sigma_{\vartheta(\tau,\sigma)}(\tau) \geq \widehat{V}^{\mathcal{N}}_{\vartheta(\tau,\sigma)}(\tau, \sigma). \tag{3.12}$$

Proof. We fix $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $\tau \in \mathcal{S}^{\mathcal{N}}$. Recall that the Snell envelope U^σ is $\mathcal{H}^\sigma$-predictable and, by Proposition 3.2, for any $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $U^\sigma(\tau)$ is a $(\mathbb{Q}, \mathbb{F})$-supermartingale. By applying the Doob decomposition theorem, we obtain the unique decomposition $U^\sigma(\tau) = M^\sigma(\tau) - A^\sigma(\tau)$ where $M^\sigma(\tau)$ is a $(\mathbb{Q}, \mathbb{F})$-martingale and $A^\sigma(\tau)$ is an increasing, integrable, $\mathbb{F}$-predictable process such that $A_0^\sigma(\tau) = 0$.

Let us take arbitrary $\tau, \tau' \in \mathcal{S}^{\mathcal{N}}$ and let us define the $\mathbb{F}$-stopping time

$$\rho = \rho(\tau, \tau') := \inf\left\{u \in [0, T] : h_u(\tau, \sigma) \neq h_u(\tau', \sigma)\right\} \wedge T. \tag{3.13}$$

Since the mapping U^σ is $\mathcal{H}^\sigma$-predictable the equality $U^\sigma(\tau) = U^\sigma(\tau')$ holds on $[\![0, \rho]\!]$, meaning that $U^\sigma_{\cdot\wedge\rho}(\tau) = U^\sigma_{\cdot\wedge\rho}(\tau')$. By the uniqueness of the Doob decomposition, we thus obtain

$$M^\sigma_{\cdot\wedge\rho}(\tau) = M^\sigma_{\cdot\wedge\rho}(\tau'), \quad A^\sigma_{\cdot\wedge\rho}(\tau) = A^\sigma_{\cdot\wedge\rho}(\tau'), \tag{3.14}$$

For a fixed t, we define the process $\widehat{M}^\sigma_u(\tau) := M^\sigma_u(\tau) - A^\sigma_t(\tau)$, $u \in [t,T]$, which is a $(\mathbb{Q},\mathbb{F})$-martingale. Then we deduce from (3.14) that $\widehat{M}^\sigma(\tau) = \widehat{M}^\sigma(\tau')$ on $[\![t,\rho]\!]$. We conclude that the mapping $\widehat{M}^\sigma : \mathcal{H}^\sigma \times [t,T] \times \Omega \to \mathbb{R}$ is $\mathcal{H}^\sigma$-predictable and thus it satisfies the assumptions of Lemma 3.6 on $[t,T]$ (see also Remark 3.2).

Let $\phi^\sigma \in \Phi_t(\mathcal{H}^\sigma)$ be the mapping given by Lemma 3.6 and Remark 3.2, for the mapping $\widehat{M}^\sigma$. To establish (3.11) and (3.12), we observe that, for all $u \in [t,T]$,

$$\widehat{Z}_u(\phi^\sigma(\tau)) = \widehat{M}^\sigma_u(\tau) = M^\sigma_u(\tau) - A^\sigma_t(\tau) = U^\sigma_u(\tau) + A^\sigma_u(\tau) - A^\sigma_t(\tau) \geq U^\sigma_u(\tau),$$

where the inequality becomes equality when $u = t$. This shows that all inequalities in (3.12) are valid. □

3.4.1 No-Arbitrage Bounds in Discrete Time

Proposition 3.4 has the following interpretation. Suppose that the issuer also holds the tranche $-\mathcal{N}$ and chooses the strategy σ. Assume also we are at time t and the holder of the tranche $\mathcal{N}$ played a strategy τ on $[0,t-1]$. Then by implementing the $\mathcal{H}^\sigma$-adapted trading strategy $\phi^\sigma(\tau)$, the issuer is super-hedging his position for all $u \in [t,T]$, so that the wealth $Z_u(\phi^\sigma(\tau))$ of his portfolio will always cover the required payoff for the holder of the tranche $\mathcal{N}$, no matter how the holder of $\mathcal{N}$ will decide to play on $[t,T]$. By construction, $\phi^\sigma(\tau)$ is also the cheapest (as of time t) of such super-hedging trading strategies for a predetermined σ, and the discounted wealth at time t of the issuer's portfolio is equal to the value $U^\sigma_t(\tau)$ of the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind.

By Proposition 3.1, an issuer's arbitrage is precluded when the price of the tranche $\mathcal{N}$ is bounded above by the value of any issuer's super hedging strategy. This implies that $U^\sigma_t(\tau)$ is an upper bound of the discounted arbitrage price at time t. In Theorem 3.1, we show that the best upper bound is obtained by taking the infimum of $U^\sigma_t(\tau)$ over all possible choices of $\sigma \in \mathcal{S}^{-\mathcal{N}}_t$.

Theorem 3.1. *Let us fix $\mathcal{N} \subseteq \mathcal{M}$ and let us consider the tranche $\mathrm{MGC}^{\mathcal{N}}_t(\mathcal{G},s)$. There is no arbitrage in $(B,S,\mathcal{N})$ at time t if and only if the discounted price $\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) = B_t^{-1}\pi^{\mathcal{N}}_t(\mathcal{G},s)$ satisfies*

$$\operatorname*{ess\,sup}_{\tau\in\mathcal{S}^{\mathcal{N}}_t(s)} \operatorname*{ess\,inf}_{\sigma\in\mathcal{S}^{-\mathcal{N}}_t(s)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma)\,\Big|\,\mathcal{F}_t\Big) \leq \widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) \leq \operatorname*{ess\,inf}_{\sigma\in\mathcal{S}^{-\mathcal{N}}_t(s)} \operatorname*{ess\,sup}_{\tau\in\mathcal{S}^{\mathcal{N}}_t(s)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma)\,\Big|\,\mathcal{F}_t\Big)$$

where $\mathbb{Q}$ is the unique martingale measure for the market model (B,S).

Proof. We will prove that the upper bound holds if and only if no issuer's arbitrage exists. The statement regarding the lower bound can be established by applying analogous arguments to the game contingent claim $\mathrm{MGC}(\widetilde{\mathcal{G}})$ with the payoffs defined by $\widetilde{V}^{-\mathcal{N}}(\tau,\sigma) = -V^{\mathcal{N}}(\tau,\sigma)$. Note that the upper bound can be rewritten as

$$\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) \leq \operatorname*{ess\,inf}_{\sigma\in\mathcal{S}^{-\mathcal{N}}_t(s)} U^\sigma_t(\tau). \tag{3.15}$$

First step. We first show that if there is no issuer's arbitrage at time t then (3.15) holds. From Proposition 3.4, we deduce that for any fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$, there exists an issuer's super-hedging strategy $(\phi^\sigma(\tau),\sigma)$ on $[t,T]$ with the discounted value at time t equal to $U^\sigma_t(\tau)$. Therefore, by Proposition 3.1, if no issuer's arbitrage exists at time t, then $\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) \leq U^\sigma_t(\tau)$. The upper bound (3.15) follows, as we may apply this argument to every $\sigma \in \mathcal{S}^{-\mathcal{N}}_t(s)$.

Second step. We will now show that if an issuer's arbitrage exists at time t then (3.15) fails to hold. Assume that (ϕ^σ,σ) is an issuer's arbitrage, in the sense of Definition 3.1. Then there exists an event $A \in \mathcal{F}_t$ with a positive probability $\mathbb{P}(A)$ (or, equivalently, $\mathbb{Q}(A)$) such that, on the event A,

$$\widehat{Z}_t(\phi^\sigma) < \widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s), \quad \widehat{Z}_{\vartheta(s)}(\phi^\sigma) \geq \widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma), \quad \forall\, \tau \in \mathcal{S}^{\mathcal{N}}_t(s).$$

Since $\widehat{Z}(\phi^\sigma)$ is a $(\mathbb{Q},\mathbb{F})$-martingale, from the optional sampling theorem we obtain, on the event A,

$$\widehat{Z}_t(\phi^\sigma) = \mathbb{E}_{\mathbb{Q}}\Big(\widehat{Z}_{\vartheta(s)}(\phi^\sigma)\,\big|\,\mathcal{F}_t\Big) \geq \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma)\,\big|\,\mathcal{F}_t\Big), \quad \forall\,\tau\in\mathcal{S}^{\mathcal{N}}_t(s).$$

Hence, by the definition of the essential supremum, we obtain on A

$$\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) > \widehat{Z}_t(\phi^\sigma) \geq \operatorname*{ess\,sup}_{\tau\in\mathcal{S}^{\mathcal{N}}_t(s)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma)\,\big|\,\mathcal{F}_t\Big) = U^\sigma_t(\tau) \geq \operatorname*{ess\,inf}_{\sigma\in\mathcal{S}^{-\mathcal{N}}_t(s)} U^\sigma_t(\tau).$$

Consequently,

$$\mathbb{Q}\Big(\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) > \operatorname*{ess\,inf}_{\sigma\in\mathcal{S}^{-\mathcal{N}}_t(s)} U^\sigma_t(\tau)\Big) > 0$$

and thus (3.15) fails to hold. □

3.5 Continuous-Time Case

In this section, we consider an arbitrage-free and complete market model with the unique martingale measure $\mathbb{Q}$, for instance, the multi-dimensional Black-Scholes model (see Karatzas and Shreve [23]). We emphasise that, in the continuous-time setup, the filtration $\mathbb{F}$ is assumed to satisfy the usual conditions of $\mathbb{Q}$-completeness and right-continuity. For $i=1,\ldots,d$, the discounted stock price processes $\widehat{S}^1,\ldots,\widehat{S}^d$ are $(\mathbb{Q},\mathbb{F})$-local martingales. Let us denote by $\mathcal{L}(\widehat{S})$ the class of all $\mathbb{R}^d$-valued, $\mathbb{F}$-predictable processes that are $\widehat{S}$-integrable on $[0,T]$. More precisely, we identify any processes ϕ and $\widetilde{\phi}$ from $\mathcal{L}(\widehat{S})$ such that

$$\int_{(0,t]} \phi_u\,d\widehat{S}_u = \int_{(0,t]} \widetilde{\phi}_u\,d\widehat{S}_u, \quad \forall\,t\in[0,T],$$

and we write $\phi \overset{\widehat{S}}{=} \widetilde{\phi}$ in that case. Due to the local properties of the Itô stochastic integral, for any stopping time $\rho\in\mathcal{T}_{[0,T]}$ and any process $\phi\in\mathcal{L}(\widehat{S})$, if $\int_{(0,\cdot]}\phi_u\,d\widehat{S}_u = 0$ on $[\![0,\rho]\!]$, meaning that

$$\int_{(0,t\wedge\rho]} \phi_u\,d\widehat{S}_u = 0, \quad \forall\,t\in[0,T],$$

then $\phi \overset{\widehat{S}}{=} 0$ on $[\![0,\rho]\!]$. When searching for super-hedging strategies, we need to modify accordingly Definition 2.2 of an $\mathcal{H}$-predictable mapping.

Definition 3.5. A mapping $\phi\colon\mathcal{H}\times[0,T]\times\Omega\to\mathbb{R}^d$ is $(\widehat{S},\mathcal{H})$-*predictable* if $\phi(h)\in\mathcal{L}(\widehat{S})$ for every $h\in\mathcal{H}$ and, for every $h,h'\in\mathcal{H}$, the equality $\phi(h)\overset{\widehat{S}}{=}\phi(h')$ holds on $[\![0,\rho]\!]$ where $\rho=\rho(h,h')$ is given by (2.1).

Throughout Section 3.5, we work under the following standing assumption, which complements Assumption 3.1.

Assumption 3.2. The discounted stock price $\widehat{S}=(\widehat{S}^1,\ldots,\widehat{S}^d)$ has the *predictable representation property* with respect to the filtration $\mathbb{F}$, that is, for any $(\mathbb{Q},\mathbb{F})$-martingale M there exists a process $\phi\in\mathcal{L}(\widehat{S})$ such that $M=\int_{(0,\cdot]}\phi_u\,d\widehat{S}_u$.

3.5.1 Right-Continuous Games

Recall that the $\mathcal{H}^\sigma$-predictable Snell envelope $U^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale for any fixed $(\tau,\sigma)\in\mathcal{S}$. To be able to apply the Doob-Meyer decomposition theorem, it suffices to show that $U^\sigma(\tau)$ is an RCLL supermartingale of class (D). In view of part (ii) in Proposition 3.2, it thus remains to show that $U^\sigma(\tau)$ admits an RCLL modification. To this end, we recall the classic Doob regularisation theorem (see, e.g., Theorem 6.27 in Kallenberg [19]).

Theorem 3.2. *A $(\mathbb{Q},\mathbb{F})$-supermartingale U admits an RCLL modification if and only if the function $\mathbb{E}_{\mathbb{Q}}(U_t)$, $t\in[0,T]$ is right-continuous.*

By applying this result to our case, we see that the supermartingale $U^\sigma(\tau)$ has an RCLL modification whenever the expected value $\mathbb{E}_{\mathbb{Q}}(U^\sigma_t(\tau))$ is right-continuous. Note also that, from the proof of part (i) in Lemma 3.4, we deduce that

$$\mathbb{E}_{\mathbb{Q}}(U^\sigma_t(\tau)) = \sup_{\tau'\in\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big). \tag{3.16}$$

The right-continuity of the function given by (3.16) is not guaranteed, in general. Therefore, we need to impose additional regularity conditions on the game $\mathcal{G}$ under the martingale measure $\mathbb{Q}$. To be more specific, in this subsection, we work under the following assumption.

Assumption 3.3. A continuous-time game $\mathcal{G}$ satisfies the following *right-continuity condition*: for any $\mathcal{N}\subseteq\mathcal{M}$ and arbitrary $s=(\tau,\sigma)\in\mathcal{S}^{-\mathcal{N}}\times\mathcal{S}^{\mathcal{N}}$, the following equalities hold for every $t\in[0,T]$,

$$\lim_{u\downarrow t}\sup_{\tau'\in\mathcal{S}^{\mathcal{N}}_u(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big) = \sup_{\tau'\in\mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big), \tag{3.17}$$

$$\lim_{u\downarrow t}\inf_{\sigma'\in\mathcal{S}^{-\mathcal{N}}_u(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau,\sigma')}(\tau,\sigma')\Big) = \inf_{\sigma'\in\mathcal{S}^{-\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau,\sigma')}(\tau,\sigma')\Big). \tag{3.18}$$

Any game $\mathcal{G}$ satisfying Assumption 3.3 is referred to as a *right-continuous game*. The motivation for this name stems from the fact that Assumption 3.3 allows us to establish the existence of an RCLL modification of the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind, as the following result shows.

Corollary 3.1. *Assume that a game $\mathcal{G}$ is right-continuous. Then, for any fixed $(\tau,\sigma)\in\mathcal{S}$, the process $U^\sigma(\tau)$ has an RCLL modification which is a $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D).*

Proof. From Assumption 3.3, it follows that the function given by the right-hand side in (3.16) is right-continuous. Hence the assertion is an immediate consequence of Proposition 3.2 and Theorem 3.2. □

Lemma 3.7. *Suppose that the mapping $\widehat{M}:\mathcal{H}\times[0,T]\times\Omega\to\mathbb{R}$ satisfies the following properties:*
(i) for any $s\in\mathcal{S}$, the process $\widehat{M}(s)=\widehat{M}(h(s))$ is an RCLL $(\mathbb{Q},\mathbb{F})$-martingale,
(ii) $\widehat{M}$ is $\mathcal{H}$-predictable.
Then there exists an $(\widehat{S},\mathcal{H})$-predictable mapping ϕ such that, for every $s\in\mathcal{S}$, the discounted wealth of a trading strategy $\phi(s)$ satisfies $\widehat{Z}(\phi(s))=\widehat{M}(s)$. The statement is still valid if we replace $\mathcal{H}$ and $[0,T]$ by $\mathcal{H}^\sigma$ and $[t,T]$, respectively.

Proof. For any fixed $s\in\mathcal{S}$, by the postulated predictable representation property of $\widehat{S}$ with respect to $\mathbb{F}$ (see Assumption 3.2), there exists an $\mathbb{F}$-predictable, $\mathbb{R}^d$-valued, integrable process $\phi(s)$ such that

$$\widehat{M}_t(s)=\widehat{M}_0(s)+\sum_{j=1}^d\int_0^t\phi^j_u(s)\,d\widehat{S}^j_u,\quad t\in[0,T].$$

It is routine to check that $\phi(s)$ is indeed an admissible trading strategy with $\widehat{Z}(\phi(s))=\widehat{M}(s)$. For any $s,s'\in\mathcal{S}$, we define the random time (see (2.1))

$$\rho(s,s'):=\inf\big\{t\in[0,T]: h_t(s)\neq h_t(s')\big\}\wedge T. \tag{3.19}$$

Then $\rho(s,s')$ is an $\mathbb{F}$-stopping time and, for every $t\in[0,T]$, the event $\{h_{[0,t)}(s)=h_{[0,t)}(s')\}$ coincides with $\{t\leq\rho(s,s')\}$. By the $\mathcal{H}$-predictability of the mapping $\widehat{M}$, the martingale $\widehat{M}(s)-\widehat{M}(s')$ vanishes on $[\![0,\rho]\!]$, and thus, one again we conclude that $\phi(s)-\phi(s')\overset{\widehat{S}}{=}0$ on $[\![0,\rho]\!]$. This shows that the mapping ϕ is $(\widehat{S},\mathcal{H})$-predictable. □

The following proposition, which holds under Assumption 3.3, is the continuous-time analogue of Proposition 3.4. We fix here t, and thus all mappings are considered on $[t,T]$, rather than $[0,T]$.

Proposition 3.5. *Fix $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $t \in [0,T]$. There exists a mapping $\phi^\sigma : \mathcal{H}^\sigma \times [t,T] \times \Omega \to \mathbb{R}$, which is $(\widehat{S}, \mathcal{H}^\sigma)$-predictable on $[t,T]$, such that for every $\tau \in \mathcal{S}^{\mathcal{N}}$, on the event $\{\vartheta(\tau,\sigma) \geq t\}$, the pair $(\phi^\sigma(\tau), \sigma)$ is an issuer's super-hedging strategy and its discounted wealth at time t is equal to the value $U_t^\sigma(\tau)$ of the $\mathcal{H}^\sigma$-predictable Snell envelope of the first kind, so that*

$$\widehat{Z}_t(\phi^\sigma(\tau)) = U_t^\sigma(\tau) = \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma) \,\Big|\, \mathcal{F}_t\Big), \tag{3.20}$$

and

$$\widehat{Z}_{\vartheta(s)}(\phi^\sigma(\tau)) \geq U^\sigma_{\vartheta(s)}(\tau) \geq \widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma). \tag{3.21}$$

Proof. The proof is analogous to the proof of Proposition 3.4. From Proposition 3.1, we know that, for any $\tau \in \mathcal{S}^{\mathcal{N}}$, the Snell envelope $U^\sigma(\tau)$ is an RCLL $\mathbb{Q}$-supermartingale of class (D). In view of the Doob-Meyer decomposition theorem, there exists a continuous, uniformly integrable $(\mathbb{Q},\mathbb{F})$-martingale $M^\sigma(\tau)$ and an increasing, $\mathbb{F}$-predictable process $A^\sigma(\tau)$ satisfying $U^\sigma(\tau) = M^\sigma(\tau) - A^\sigma(\tau)$, $A^\sigma_0(\tau) = 0$ and this decomposition is unique. We define a candidate for the discounted wealth process by setting $\widehat{M}^\sigma_u(\tau) = M^\sigma_u(\tau) - A^\sigma_t(\tau)$ for all $u \in [t,T]$. It clear that, for any fixed $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $\widehat{Z}^\sigma(\tau)$ is a uniformly integrable $(\mathbb{Q},\mathbb{F})$-martingale.

We will show that the mapping $\widehat{M}^\sigma$ is $\mathcal{H}^\sigma$-predictable on $[t,T]$. To this end, we take arbitrary $s = (\tau,\sigma)$, $s' = (\tau',\sigma)$ and we define the $\mathbb{F}$-stopping time $\rho = \rho(\tau,\tau')$ using equation (3.13). Since U^σ is $\mathcal{H}^\sigma$-predictable, we have that $U^\sigma(\tau) = U^\sigma(\tau')$ on $[\![0,\rho]\!]$. From the uniqueness of the Doob-Meyer decomposition, we thus obtain $M^\sigma(\tau) = M^\sigma(\tau')$ and $A^\sigma(\tau) = A^\sigma(\tau')$ on $[\![0,\rho]\!]$. We deduce easily that $\widehat{M}^\sigma(\tau) = \widehat{M}^\sigma(\tau')$ on $[\![t,\rho]\!]$ and thus the $\mathcal{H}^\sigma$-predictability of $\widehat{M}^\sigma$ on $[t,T]$ is established. From Lemma 3.7, we deduce the existence of a mapping $\phi^\sigma : \mathcal{H}^\sigma \times [t,T] \times \Omega \to \mathbb{R}$, which is $(\widehat{S}, \mathcal{H}^\sigma)$-predictable and such that the equality $\widehat{Z}(\phi^\sigma(\tau)) = \widehat{M}^\sigma(\tau)$ holds on $[t,T]$.

It now remains to check that (3.20) and (3.21) are valid. For any $\mathbb{F}$-stopping time $\vartheta(s)$ taking values in $[t,T]$, we obtain, on the event $\{\vartheta(s) \geq t\}$,

$$\widehat{Z}_{\vartheta(s)}(\phi^\sigma(\tau)) = \widehat{M}_{\vartheta(s)}(\phi^\sigma(\tau)) = M^\sigma_{\vartheta(s)}(\tau) - A^\sigma_t(\tau) = U^\sigma_{\vartheta(s)}(\tau) + A^\sigma_{\vartheta(s)}(\tau) - A^\sigma_t(\tau) \geq U^\sigma_{\vartheta(s)}(\tau)$$

with equality holding on the event $\{\vartheta(s) = t\}$. This completes the proof of the theorem. □

3.5.2 General Games

Assumption 3.3 is in fact too strong, since it is easy to produce examples of games for which it fails to hold. Therefore, this assumption will be now relaxed, and we will show that a counterpart of Corollary 3.1 can still be established (see Corollary 3.2). To this end, we will introduce an alternative definition of the $\mathcal{H}^\sigma$-predictable Snell envelope (see Definition 3.6), so that we need first to analyze some basic relationships between the Snell envelopes U^σ and $\widetilde{U}^\sigma$. From Section 3.3, we already know that the processes $U^\sigma(\tau)$ and $\widetilde{U}^\sigma(\tau)$ are $(\mathbb{Q},\mathbb{F})$-supermartingales, and thus their expected values are decreasing functions of time.

Lemma 3.8. *(i) For any $\tau \in \mathcal{S}^{\mathcal{N}}$ and $t < u$, we have*

$$U^\sigma_t(\tau) \geq \widetilde{U}^\sigma_t(\tau) \geq \mathbb{E}_{\mathbb{Q}}\big(U^\sigma_u(\tau) \,\big|\, \mathcal{F}_t\big) \geq \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_u(\tau) \,\big|\, \mathcal{F}_t\big). \tag{3.22}$$

(ii) The following set $b^\sigma(\tau)$ is countable

$$b^\sigma(\tau) := \big\{t \in [0,T] : \mathbb{E}_{\mathbb{Q}}(U^\sigma_t(\tau)) > \mathbb{E}_{\mathbb{Q}}(\widetilde{U}^\sigma_t(\tau))\big\}.$$

(iii) If $t \notin b^\sigma(\tau)$, then $U^\sigma_t(\tau) = \widetilde{U}^\sigma_t(\tau)$.

Proof. (i) The asserted inequalities follow from part (i) in Lemma 3.4 and the following obvious inclusions $\widetilde{\mathcal{S}}_u^{\mathcal{N}}(\tau,\sigma) \subseteq \mathcal{S}_u^{\mathcal{N}}(\tau,\sigma) \subseteq \widetilde{\mathcal{S}}_t^{\mathcal{N}}(\tau,\sigma) \subseteq \mathcal{S}_t^{\mathcal{N}}(\tau,\sigma)$.

(ii) We start by noting that the inequality $\mathbb{E}_{\mathbb{Q}}(U_t^\sigma(\tau)) \geq \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_t^\sigma(\tau))$ holds for all t. For any (τ,σ), we define the set $b^\sigma(\tau)$ of dates for which the inequality is strict. In view of (3.22), the decreasing functions $\mathbb{E}_{\mathbb{Q}}(U_t^\sigma(\tau))$ and $\mathbb{E}_{\mathbb{Q}}(\widetilde{U}_t^\sigma(\tau))$ have the same right-hand and left-hand limits for every t. Therefore, they are equal at any t for which at least one of them (and thus both) are continuous. To conclude the proof of part (ii), we observe that the set of points of discontinuity of a decreasing function is countable.

(iii) We observe that if $t \notin b^\sigma(\tau)$ then $\mathbb{E}_{\mathbb{Q}}(U_t^\sigma(\tau)) = \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_t^\sigma(\tau))$. Hence the assertion is a direct consequence of the inequality $U_t^\sigma(\tau) \geq \widetilde{U}_t^\sigma(\tau)$ (see part (i)). □

We define the mappings B^σ and $\widetilde{B}^\sigma$, which capture the cumulative differences between U^σ and $\widetilde{U}^\sigma$, by setting, for all $t \in [0,T]$ and $\tau \in \mathcal{S}^{\mathcal{N}}$,

$$B_t^\sigma(\tau) := \sum_{u \in b_t^\sigma(\tau)} U_u^\sigma(\tau) - \widetilde{U}_u^\sigma(\tau),$$
$$\widetilde{B}_t^\sigma(\tau) := \sum_{u \in \widetilde{b}_t^\sigma(\tau)} U_u^\sigma(\tau) - \widetilde{U}_u^\sigma(\tau),$$

where we denote $b_t^\sigma(\tau) = [0,t) \cap b^\sigma(\tau)$ and $\widetilde{b}_t^\sigma(\tau) = [0,t] \cap b^\sigma(\tau)$ and we apply the usual convention that the empty sum equals zero. In view of part (iii) in Lemma 3.8, we have $B_t^\sigma(\tau) = \widetilde{B}_t^\sigma(\tau)$ for every $t \notin b^\sigma(\tau)$.

Lemma 3.9. *(i) The processes $B^\sigma(\tau)$ and $\widetilde{B}^\sigma(\tau)$ are non-negative and increasing.*
(ii) The mapping B^σ ($\widetilde{B}^\sigma$, resp.) is $\mathcal{H}^\sigma$-predictable ($\mathcal{H}^\sigma$-adapted, resp.).
(iii) For any $t \leq u < v \leq T$, we have the inequality

$$\mathbb{E}_{\mathbb{Q}}\Big(\widetilde{B}_v^\sigma(\tau) - \widetilde{B}_u^\sigma(\tau) \,\Big|\, \mathcal{F}_t\Big) \leq \mathbb{E}_{\mathbb{Q}}\Big(\widetilde{U}_u^\sigma(\tau) - \widetilde{U}_v^\sigma(\tau) \,\Big|\, \mathcal{F}_t\Big).$$

Proof. (i) By part (ii) in Lemma 3.8, the set $b^\sigma(\tau)$ is countable and thus the inequality $U_t^\sigma(\tau) \geq \widetilde{U}_t^\sigma(\tau)$ holds for all $t \in b^\sigma(\tau)$, with probability 1. Therefore, $B^\sigma(\tau)$ and $\widetilde{B}^\sigma(\tau)$ are non-negative increasing processes.

(ii) To show that B^σ is $\mathcal{H}^\sigma$-predictable, it suffices to check that $B^\sigma(\tau) = B^\sigma(\tau')$ on $[\![0,\rho]\!]$ where $\rho = \rho(\tau,\tau')$ is given by (3.13). Since, by part (iii) in Lemma 3.8, equality $U_t^\sigma(\tau) = \widetilde{U}_t^\sigma(\tau)$ holds for every $t \notin b^\sigma(\tau)$, we obtain

$$B_t^\sigma(\tau) = \sum_{u \in b_t^\sigma(\tau) \cup b_t^\sigma(\tau')} U_u^\sigma(\tau) - \widetilde{U}_u^\sigma(\tau),$$
$$B_t^\sigma(\tau') = \sum_{u \in b_t^\sigma(\tau) \cup b_t^\sigma(\tau')} U_u^\sigma(\tau') - \widetilde{U}_u^\sigma(\tau').$$

Both sums are now over same countable index set, and thus the $\mathcal{H}^\sigma$-predictability of B^σ follows from the definition of $b_t^\sigma(\tau)$ and $b_t^\sigma(\tau')$, since the $\mathcal{H}^\sigma$-adaptedness of U^σ and $\widetilde{U}^\sigma$ implies that $U^\sigma(\tau) = U^\sigma(\tau')$ and $\widetilde{U}^\sigma(\tau) = \widetilde{U}^\sigma(\tau')$ on $[\![0,\rho)\!]$. The $\mathcal{H}^\sigma$-adaptedness of $\widetilde{B}^\sigma$ can be shown using similar arguments.

(iii) From parts (i) and (ii) in Lemma 3.8, the following set of countable stochastic open intervals

$$\Big\{\Big(\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_w^\sigma(\tau) \,\big|\, \mathcal{F}_t\big),\ \mathbb{E}_{\mathbb{Q}}\big(U_w^\sigma(\tau) \,\big|\, \mathcal{F}_t\big)\Big) : w \in (u,v] \cap b^\sigma(\tau)\Big\}$$

consists of pairwise disjoint stochastic intervals lying within

$$\Big[\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_v^\sigma(\tau) \,\big|\, \mathcal{F}_t\big),\ \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_u^\sigma(\tau) \,\big|\, \mathcal{F}_t\big)\Big],$$

almost surely. Consequently, the non-negative random variable $\widetilde{B}^\sigma_v(\tau) - \widetilde{B}^\sigma_u(\tau)$ satisfies

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}}\Big(\widetilde{B}^\sigma_v(\tau) - \widetilde{B}^\sigma_u(\tau)\,\big|\,\mathcal{F}_t\Big) &= \sum_{w\in(u,v]\cap b^\sigma(\tau)} \mathbb{E}_{\mathbb{Q}}\Big(U^\sigma_w(\tau)\,\big|\,\mathcal{F}_t\Big) - \mathbb{E}_{\mathbb{Q}}\Big(\widetilde{U}^\sigma_w(\tau)\,\big|\,\mathcal{F}_t\Big)\\ &\le \mathbb{E}_{\mathbb{Q}}\Big(\widetilde{U}^\sigma_u(\tau) - \widetilde{U}^\sigma_v(\tau)\,\big|\,\mathcal{F}_t\Big)\end{aligned}$$

where the equality follows from Fubini's theorem The inequality $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{B}^\sigma_T(\tau)\big) < \infty$ now follows, and thus the process $\widetilde{B}^\sigma(\tau)$ is $\mathbb{Q}$-integrable. □

The following lemma establishes right-continuity conditions, which can be seen as an alternative to Assumption 3.3.

Lemma 3.10. *For every $\mathcal{N} \subseteq \mathcal{M}$ and $s = (\tau,\sigma) \in \mathcal{S}^{-\mathcal{N}} \times \mathcal{S}^{\mathcal{N}}$, the following equalities hold for every $t \in [0,T]$,*

$$\lim_{u\downarrow t} \sup_{\tau'\in\widetilde{\mathcal{S}}^{\mathcal{N}}_u(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big) = \sup_{\tau'\in\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big), \tag{3.23}$$

$$\lim_{u\downarrow t} \inf_{\sigma'\in\widetilde{\mathcal{S}}^{-\mathcal{N}}_u(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau,\sigma')}(\tau,\sigma')\Big) = \inf_{\sigma'\in\widetilde{\mathcal{S}}^{-\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau,\sigma')}(\tau,\sigma')\Big). \tag{3.24}$$

Proof. By symmetry, it suffices to only prove the first equality. From part (i) in Proposition 3.3, we have the equality

$$\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_t(\tau)\big) = \sup_{\tau'\in\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big),$$

so that (3.23) is equivalent to the right-continuity of the function $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_t(\tau)\big)$, $t\in[0,T]$. Since for $t<w$, we have $\widetilde{\mathcal{S}}^{\mathcal{N}}_w(\tau,\sigma) \subseteq \widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)$, and thus the function $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_t(\tau)\big)$ is decreasing. Hence the right-hand limit $\lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_w(\tau)\big)$ is well-defined and bounded above by $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_t(\tau)\big)$.

Clearly, to establish equality (3.23), that is, the right-continuity of the function $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma(\tau)\big)$, it is enough to show that $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_t(\tau)\big) \le \lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}^\sigma_w(\tau)\big)$ for all $t\in[0,T]$. For this purpose, we fix $\widetilde{\tau}\in\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)$ and we consider the stopping time ρ given by

$$\rho := \rho(\tau,\widetilde{\tau}) = \inf\big\{u\in[0,T] : h_u(\tau,\sigma) \neq h_u(\sigma,\widetilde{\tau})\big\}\wedge T.$$

Since $\widetilde{\tau}\in\widetilde{\mathcal{S}}^{\mathcal{N}}_t(\tau,\sigma)$, we must have $h_{[0,t]}(\tau,\sigma) = h_{[0,t]}(\widetilde{\tau},\sigma)$ and thus $\rho > t$. Consider the process $\big(\widetilde{U}^\sigma_w(\tau) - \widetilde{V}^{\mathcal{N}}_w(\widetilde{\tau},\sigma)\big)\mathbb{1}_{\{w<\rho\}}$, for $w\in(t,T]$, where

$$\widetilde{V}^{\mathcal{N}}_w(\widetilde{\tau},\sigma) := \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\widetilde{\tau},\sigma)}(\widetilde{\tau},\sigma)\,\Big|\,\mathcal{F}_w\Big).$$

Assumption 2.1 and the finiteness of the sets of actions yield $\{w<\rho\} = \{h_{[0,w]}(\tau,\sigma) = h_{[0,w]}(\sigma,\widetilde{\tau})\}$. Consequently, we obtain

$$\widetilde{V}^{\mathcal{N}}_w(\widetilde{\tau},\sigma)\mathbb{1}_{\{w<\rho\}} \le \operatorname*{ess\,sup}_{\tau'\in\widetilde{\mathcal{S}}^{\mathcal{N}}_w(\tau,\sigma)} \widetilde{V}^{\mathcal{N}}_w(\tau',\sigma)\mathbb{1}_{\{w<\rho\}} = \widetilde{U}^\sigma_w(\tau)\mathbb{1}_{\{w<\rho\}},$$

which in turn implies that

$$\big(\widetilde{U}^\sigma_w(\tau) - \widetilde{V}^{\mathcal{N}}_w(\widetilde{\tau},\sigma)\big)\mathbb{1}_{\{w<\rho\}} \ge 0. \tag{3.25}$$

Observe that $\widetilde{V}^{\mathcal{N}}(\widetilde{\tau},\sigma)$ is a $(\mathbb{Q},\mathbb{F})$-martingale and thus, by Theorem 3.2, it admits an RCLL modification. We also know that $\widetilde{U}^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale (see part (iii) of Proposition 3.3), but its right-continuity was not yet established. Doob's regularization theorem (see, e.g., Theorem 6.27 in Kallenberg [19]) states that for any submartingale X on $\mathbb{R}_+$ with restriction Y on the set on

positive rational numbers $\mathcal{Q}_+$, the process $Y_t^+ := \lim_{w\downarrow t, w\in\mathcal{Q}} Y_w$ is an RCLL process outside some fixed $\mathbb{Q}$-null set. By slightly modifying the proof of this theorem we may and do assume, without loss of generality, that the countable set

$$I(\tau,\sigma) := \big\{t\in[0,T] : \mathbb{E}_{\mathbb{Q}}(U_t^\sigma(\tau)) > \lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}(\widetilde{U}_w^\sigma(\tau))\big\},$$

is included in the countable, dense subset $\widetilde{\mathcal{Q}}_+$ of $\mathbb{R}_+$. So we may define, for all $t\in\widetilde{\mathcal{Q}}_+$,

$$\widetilde{U}_{t+}^\sigma(\tau) := \lim_{w\downarrow t, w\in\widetilde{\mathcal{Q}}_+} \widetilde{U}_w^\sigma(\tau).$$

One the one hand, the dominated convergence theorem (recall that we work under Assumption 3.1) yields, for all $t\in\widetilde{\mathcal{Q}}_+$,

$$\lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_w^\sigma(\tau)\big) = \lim_{w\downarrow t, w\in\widetilde{\mathcal{Q}}_+}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_w^\sigma(\tau)\big) = \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_{t+}^\sigma(\tau)\big). \tag{3.26}$$

On the other hand, by applying the dominated convergence theorem to (3.25), we obtain

$$0 \le \lim_{w\downarrow t, w\in\widetilde{\mathcal{Q}}_+}\mathbb{E}_{\mathbb{Q}}\Big(\big(\widetilde{U}_w^\sigma(\tau) - \widetilde{V}_w^{\mathcal{N}}(\widetilde{\tau},\sigma)\big)\mathbb{1}_{\{w<\rho\}}\Big) = \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_{t+}^\sigma(\tau) - \widetilde{V}_t^{\mathcal{N}}(\widetilde{\tau},\sigma)\big) \tag{3.27}$$

since the process $\mathbb{1}_{\{w<\rho\}}$, $w\in[t,T]$, is bounded, with right-continuous sample paths, and $\mathbb{1}_{\{t<\rho\}}=1$. By combining (3.26) with (3.27), we conclude that

$$\lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_w^\sigma(\tau)\big) = \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_{t+}^\sigma(\tau)\big) \ge \mathbb{E}_{\mathbb{Q}}\big(\widetilde{V}_t^{\mathcal{N}}(\widetilde{\tau},\sigma)\big) = \mathbb{E}_{\mathbb{Q}}\big(\widehat{V}^{\mathcal{N}}_{\vartheta(\widetilde{\tau},\sigma)}(\widetilde{\tau},\sigma)\big). \tag{3.28}$$

Since (3.28) holds for all $\widetilde{\tau}\in\widetilde{\mathcal{S}}_t^{\mathcal{N}}(\tau,\sigma)$, we must have, for all $t\in\widetilde{\mathcal{Q}}_+$,

$$\lim_{w\downarrow t}\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_w^\sigma(\tau)\big) \ge \sup_{\tau'\in\widetilde{\mathcal{S}}_t^{\mathcal{N}}(\tau,\sigma)}\mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\Big) = \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_t^\sigma(\tau)\big).$$

In view of the definition of the set $I(\tau,\sigma)\subset\widetilde{\mathcal{Q}}_+$, this completes the proof of the lemma. □

Corollary 3.2. *For any $(\tau,\sigma)\in\mathcal{S}^{-\mathcal{N}}\times\mathcal{S}^{\mathcal{N}}$, we have that*
(i) the process $\widetilde{U}^\sigma(\tau)$ has an RCLL modification, which is a $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D).
(ii) the process $\widetilde{B}^\sigma(\tau)$ has an RCLL modification, which is a $(\mathbb{Q},\mathbb{F})$-submartingale of class (D).

Proof. (i) In view of part (iii) in Proposition 3.3, the process $\widetilde{U}^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale. Moreover, by equality (3.23) in Lemma 3.10, its expected value $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_t^\sigma(\tau)\big)$ is a right-continuous function. The existence of an RCLL modification of $\widetilde{U}^\sigma(\tau)$ thus follows from Theorem 3.2. To show that it is $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D), we use part (iii) in Proposition 3.3 and we argue as is the proof of part (ii) in Proposition 3.2.

(ii) To show that $\widetilde{B}^\sigma(\tau)$ has an RCLL modification, it suffices to check that the function $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{B}_t^\sigma(\tau)\big)$ is right-continuous or, more explicitly,

$$\lim_{u\downarrow t}\sum_{v\in(t,u]\cap b^\sigma(\tau)}\mathbb{E}_{\mathbb{Q}}(U_v^\sigma(\tau)) - \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_v^\sigma(\tau)) = 0.$$

Since the functions $\mathbb{E}_{\mathbb{Q}}\big(U_t^\sigma(\tau)\big) \ge \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_t^\sigma(\tau)\big)$ are decreasing and the equality $\mathbb{E}_{\mathbb{Q}}\big(U_t^\sigma(\tau)\big) = \mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_t^\sigma(\tau)\big)$ holds for all $t\notin b^\sigma(\tau)$, we obtain

$$0 \le \sum_{v\in(t,u]\cap b^\sigma(\tau)}\mathbb{E}_{\mathbb{Q}}(U_v^\sigma(\tau)) - \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_v^\sigma(\tau)) \le \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_t^\sigma(\tau)) - \mathbb{E}_{\mathbb{Q}}(\widetilde{U}_u^\sigma(\tau)),$$

where the difference tends to 0 as $u\downarrow t$ since, from part (i), the function $\mathbb{E}_{\mathbb{Q}}\big(\widetilde{U}_t^\sigma(\tau)\big)$ is right-continuous. □

3.5.3 $\mathcal{H}^\sigma$-Predictable Snell Envelope of the Second Kind

We are in a position to introduce another version of the concept of the $\mathcal{H}^\sigma$-predictable Snell envelope, which will serve as a convenient tool when searching for the issuer's super-hedging strategy.

Definition 3.6. For a combined tranche $\mathrm{MGC}^{\mathcal{N}}(\mathcal{G})$ and a fixed $\sigma \in \mathcal{S}^{-\mathcal{N}}$, the *$\mathcal{H}^\sigma$-predictable Snell envelope of the second kind* $\bar{U}^\sigma : \mathcal{H}^\sigma \times [0,T] \times \Omega \to \mathbb{R}$ is given by

$$\bar{U}^\sigma_t(\tau) := \widetilde{U}^\sigma_t(\tau) + \widetilde{B}^\sigma_t(\tau) = U^\sigma_t(\tau) + B^\sigma_t(\tau) \tag{3.29}$$

where the second equality follows from part (iii) in Lemma 3.8.

Intuitively, the difference $\widetilde{B}^\sigma_t(\tau) = \bar{U}^\sigma_t(\tau) - \widetilde{U}^\sigma_t(\tau)$ represents the cumulative adjustments to the process $\widetilde{U}^\sigma(\tau)$ that are needed to ensure that the inequality $\bar{U}^\sigma_t(\tau) \geq U^\sigma_t(\tau)$ holds for all t. It is fair to say, however, that $\bar{U}^\sigma(\tau)$ does not enjoy as nice interpretation as $U^\sigma(\tau)$, so it should merely be seen as a technical tool to derive the desired results, Propositions 3.6 and 3.7.

Proposition 3.6. *For any $\sigma \in \mathcal{S}^{-\mathcal{N}}$:*
(i) the mapping $\bar{U}^\sigma$ is an $\mathcal{H}^\sigma$-predictable,
(ii) for any $\tau \in \mathcal{S}^{\mathcal{N}}$, the process $\bar{U}^\sigma(\tau)$ is an RCLL $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D).

Proof. (i) It suffices to observe that the second equality in (3.29) yields $\bar{U}^\sigma = U^\sigma + B^\sigma$ where U^σ and B^σ are $\mathcal{H}^\sigma$-predictable by Lemmas 3.1 and 3.9, respectively.

(iii) The first equality in (3.29) reads $\bar{U}^\sigma(\tau) = \widetilde{U}^\sigma(\tau) + \widetilde{B}^\sigma(\tau)$. Hence, from part (ii) in Lemma 3.9, the process $\bar{U}^\sigma(\tau)$ is a $(\mathbb{Q},\mathbb{F})$-supermartingale. Next, from Corollary 3.2, it follows that $\bar{U}^\sigma(\tau)$ has an RCLL modification and is of class (D). □

We will now prove the existence of a super-hedging strategy on $[t,T]$. It is worth noting that we would like to employ here the following arguments, for any fixed t:
(i) the process $\bar{U}^\sigma_u(\tau)$, $u \in [t,T]$ is an RCLL $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D),
(ii) we may and do assume, without loss of generality, that the equality $\bar{U}^\sigma_t(\tau) = U^\sigma_t(\tau)$ holds, for instance, by replacing the process $\widetilde{B}^\sigma(\tau)$ by $\widetilde{B}^\sigma(\tau)\mathbb{1}_{[t,T]}$ in formula (3.29), that is, by neglecting the past adjustments to $\bar{U}^\sigma(\tau)$ on $[0,t)$.

Formally, in view of the second equality in (3.29), for a fixed $t \in [0,T]$, to make use of the above features (i) and (ii), it is enough to define the auxiliary process $\check{U}^\sigma_u(\tau) = \bar{U}^\sigma_u(\tau) - B_t$, $u \in [t,T]$, which obviously satisfies $\check{U}^\sigma_t(\tau) = U^\sigma_t(\tau)$ and is an RCLL $(\mathbb{Q},\mathbb{F})$-supermartingale of class (D).

Proposition 3.7. *Fix $\sigma \in \mathcal{S}^{-\mathcal{N}}$ and $t \in [0,T]$. There exists a mapping $\phi^\sigma : \mathcal{H}^\sigma \times [t,T] \times \Omega \to \mathbb{R}$, which is $(\widehat{S},\mathcal{H}^\sigma)$-predictable on $[t,T]$, such that for every $\tau \in \mathcal{S}^{\mathcal{N}}$, on the event $\{\vartheta(\tau,\sigma) \geq t\}$, the pair $(\phi^\sigma(\tau),\sigma)$ is an issuer's super-hedging strategy and its discounted wealth at time t is equal to the value $\bar{U}^\sigma_t(\tau)$ of the $\mathcal{H}^\sigma$-predictable Snell envelope of the second kind, so that*

$$\widehat{Z}_t(\phi^\sigma(\tau)) = \check{U}^\sigma_t(\tau) = U^\sigma_t(\tau) = \operatorname*{ess\,sup}_{\tau' \in \mathcal{S}^{\mathcal{N}}_t(\tau,\sigma)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}^{\mathcal{N}}_{\vartheta(\tau',\sigma)}(\tau',\sigma)\,\Big|\,\mathcal{F}_t\Big), \tag{3.30}$$

and

$$\widehat{Z}_{\vartheta(s)}(\phi^\sigma(\tau)) \geq \check{U}^\sigma_{\vartheta(s)}(\tau) \geq U^\sigma_{\vartheta(s)}(\tau) \geq \widehat{V}^{\mathcal{N}}_{\vartheta(s)}(\tau,\sigma). \tag{3.31}$$

Proof. The proof goes along the same lines as the proof of Proposition 3.5 with only minor modifications. We may now apply the Doob-Meyer decomposition $\check{U}^\sigma(\tau) = \check{M}^\sigma(\tau) - \check{A}^\sigma(\tau)$, to obtain an $\mathcal{H}^\sigma$-predictable martingale mapping $\widehat{M}^\sigma_u(\tau)$, $u \in [t,T]$, which dominates $U^\sigma(\tau)$ on $[t,T]$, and which will play the role of the wealth process of a super-hedging portfolio on $[t,T]$. It is interesting to note that we obtain the following decompositions for $(\mathbb{Q},\mathbb{F})$-supermartingales $U^\sigma(\tau)$ and $\widetilde{U}^\sigma(\tau)$

$$U^\sigma(\tau) = \check{M}^\sigma(\tau) - \check{A}^\sigma(\tau) - B^\sigma(\tau), \quad \widetilde{U}^\sigma(\tau) = \check{M}^\sigma(\tau) - \check{A}^\sigma(\tau) - \widetilde{B}^\sigma(\tau).$$

Recall that $\check{U}^\sigma_t(\tau) = U^\sigma_t(\tau)$. We then proceed as in the proof of Proposition 3.5. □

3.5.4 No-Arbitrage Bounds in Continuous Time

The following theorem furnishes the no-arbitrage bounds for a price $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ of a combined tranche in the continuous-time framework. As expected, its conclusion is exactly the same as in discrete time setup (see Theorem 3.1).

Theorem 3.3. *Assume that a continuous time game $\mathcal{G}$ satisfies Assumption 3.1. Let us fix $\mathcal{N} \subseteq \mathcal{M}$ and $t \in [0, T]$, and let us consider the tranche $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$. There is no arbitrage in $(B, S, \mathcal{N})$ if and only if the discounted price $\widehat{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) = B_t^{-1}\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ at time t satisfies*

$$\operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}_{\vartheta(s)}^{\mathcal{N}}(\tau, \sigma) \,\Big|\, \mathcal{F}_t\Big) \leq \widehat{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \leq \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \mathbb{E}_{\mathbb{Q}}\Big(\widehat{V}_{\vartheta(s)}^{\mathcal{N}}(\tau, \sigma) \,\Big|\, \mathcal{F}_t\Big). \quad (3.32)$$

Proof. In view of Propositions 3.5 and 3.7, the arguments used in the proof of Theorem 3.3 are exactly the same as the demonstration of Theorem 3.1 and thus we omit the details. □

3.6 A Counter-Example

Unlike for two-person game options, the arbitrage-free condition of Definition 3.1 restricts the price to an interval, rather than produces a single value. This is due to the fact that we are considering a general game $\mathcal{G}$ with arbitrary payoff functions. It is difficult to make further progress without being more specific about $\mathcal{G}$, for instance, by restricting attention to zero-sum games. It is worth stressing that in the non-zero case, even if the inequalities (3.32) are satisfied and arbitrage prices of single tranches are uniquely defined, they may still be problematic from the financial point of view. The following example illustrates this statement.

Example 3.1. This example is based on the famous *prisoner's dilemma.* Specifically, let $\mathcal{G}$ be a deterministic game with two players, each having the action space of $\mathcal{S}^i = \{0, 1\}$ for $i = 1, 2$. The payoff function is given by

$V(\tau, \sigma)$	$\tau = 0$	$\tau = 1$
$\sigma = 0$	$(1, 1)$	$(-1, 2)$
$\sigma = 1$	$(2, -1)$	$(0, 0)$

Note that there is an optimal equilibrium at $(\tau^*, \sigma^*) = (1, 1)$ and thus the unique value is $(0, 0)$.

Consider the contract $\mathrm{MGC}(\mathcal{G})$ in a market model with null interest rate. It is easy to see that for each individual tranches, the upper and lower prices are both \$0. Hence the arbitrage price for each tranche must be $\pi^1(\mathcal{G}) = \pi^2(\mathcal{G}) = 0$. But when we consider the collection of tranches $\mathcal{M} = \{1, 2\}$, the unique arbitrage price is

$$\pi^{\mathcal{M}}(\mathcal{G}) = \operatorname*{ess\,sup}_{s \in \mathcal{S}} \big(V^1(s) + V^2(s)\big) = V^1(1, 1) + V^2(1, 1) = 2. \quad (3.33)$$

In particular, note that $\pi^1(\mathcal{G}) + \pi^2(\mathcal{G}) \neq \pi^{\mathcal{M}}(\mathcal{G})$.

Now interpret this from the perspective of the issuer and the two prospective holders who are about to enter the contract $\mathrm{MGC}(\mathcal{G})$. For the issuer, the total selling price for the two tranches cannot be less than \$2, since there is nothing stopping the two holders from working together for a combined payoff of \$2. For the prospective holders of the tranche 1, the buying price cannot be more than \$0, since there is nothing stopping the holder of the tranche 2 from playing 1, and limiting the payoff of the tranche 1 to at most 0. In fact, it is strictly better for the holder of the tranche 2 to play 1 than 0 in all cases. A similar analysis can be made for the prospective holder of the tranche 2.

So it is impossible for any deal to strike between the three parties. The issuer will always be demanding more than what the prospective holders are willing to pay. The only way to resolve this situation is for at least two of the three parties to agree to work together and share their payoffs in some way. For example, the two holders agree to buy the two tranches for \$2 and play $(0, 0)$; or the issuer agrees to sell the tranches for \$0, but requiring one of the holders to play 1.

3.7 No-Arbitrage Versus Super-Hedging

We have shown in Theorems 3.1 and 3.3 that there is no arbitrage in the market $(B, S, \mathcal{N})$ if and only if the discounted price of the tranche $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$ satisfies

$$\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \le \widehat{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \le \overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \tag{3.34}$$

where the *upper price* $\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ and the *lower price* $\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ are defined by

$$\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) := \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma),$$
$$\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) := \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma),$$

where

$$\widehat{V}_t^{\mathcal{N}}(\tau, \sigma) := \mathbb{E}_{\mathbb{Q}}\big(\widehat{V}_\vartheta^{\mathcal{N}}(\tau, \sigma) \,\big|\, \mathcal{F}_t\big).$$

In Theorems 3.1 and 3.3, the upper and lower prices were derived using arbitrage pricing arguments. As expected, they also have an interpretation in terms of super-hedging strategies introduced in Definition 3.2. To this end, we will need the following property.

Lemma 3.11. *There exist two sequences of strategies, $(\tau_n)_{n\in\mathbb{N}}$ from $\mathcal{S}_t^{\mathcal{N}}(s)$ and $(\sigma_n)_{n\in\mathbb{N}}$ from $\mathcal{S}_t^{-\mathcal{N}}(s)$, such that the sequences*

$$\left(\operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma_n)\right)_{n\in\mathbb{N}} \quad \text{and} \quad \left(\operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma)\right)_{n\in\mathbb{N}}$$

are almost surely non-increasing and non-decreasing, respectively. Furthermore, they converge to the upper and lower prices, specifically,

$$\lim_{n\to\infty} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma_n) = \inf_{n\in\mathbb{N}} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma_n) = \overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s),$$

and

$$\lim_{n\to\infty} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma) = \sup_{n\in\mathbb{N}} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma) = \underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s).$$

Proof. Using arguments identical to Lemma 3.2, the sets

$$\left\{\operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma) : \sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)\right\}, \quad \left\{\operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma) : \tau \in \mathcal{S}_t^{\mathcal{N}}(s)\right\} \tag{3.35}$$

have the lattice property. Hence the existence of the desired sequences follows by part (i) in Lemma 3.3. □

Suppose the issuer holds the tranche $-\mathcal{N}$ and plays the strategy $\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)$ while super-hedging using (ϕ^σ, σ). The upper price $\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ is the infimum of $\widehat{Z}_t(\phi^\sigma)$ (the discounted time t value) over all possible $\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)$. Even though the infimum is not necessarily achieved, Lemma 3.11 shows that there exists a sequence of $(\sigma_n)_{n\in\mathbb{N}}$ such that $\widehat{Z}_t(\phi^{\sigma_n})$ is non-increasing and converging to $\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$. Intuitively, the buyer of the tranche $\mathcal{N}$ will be reluctant to pay more than $\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ at time t. Otherwise, his expected discounted payoff could be restricted to be lower than the price paid and, as shown by Theorems 3.1 and 3.3, leads to an issuer's arbitrage.

Similarly, the holder of the tranche $\mathcal{N}$ plays the strategy $\tau \in \mathcal{S}_t^{\mathcal{N}}(s)$ and super-hedges the negative of his payoff using (ϕ^τ, τ). The lower price $\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ is the supremum of $-\widehat{Z}_t(\phi^\tau)$ over all possible $\tau \in \mathcal{S}_t^{\mathcal{N}}(s)$. Again the supremum is not necessarily achieved, but Lemma 3.11 shows that there exists a sequence of $(\tau_n)_{n\in\mathbb{N}}$ such that $-\widehat{Z}_t(\phi^{\sigma_n})$ is non-decreasing and converges to $\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$. Intuitively, it is undesirable for the issuer to sell the tranche $\mathcal{N}$ for a price cheaper than $\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$. Otherwise the holder of the tranche $\mathcal{N}$ will be able to guarantee an expected discounted payoff greater than the price. As shown in Theorems 3.1 and 3.3, this leads to a holder's arbitrage.

4 Pricing of Multi-Person Game Claims

In the previous section, we focussed on arbitrage valuation of a fixed combined tranche $\mathcal{N}$ of $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$. In Theorems 3.1 and 3.3, we have shown that, in the absence of arbitrage in the market model $(B, S, \mathcal{N})$ (see Definition 3.1), an arbitrage price at time t of a combined tranche $\mathcal{N}$ lies between the upper and lower prices, that is, (3.34) holds. If there is no arbitrage in $(B, S, \mathcal{N})$ for *any* non-empty subset of tranches $\mathcal{N} \subseteq \mathcal{M}$, then (3.34) actually describes a total of $2(2^m - 1)$ inequalities, which are not easy to handle, in general. In this section, we shall further explore how the prices of different collections of tranches are related to one another. In particular, we address in this section the following questions:

- Apart from (3.34), should there be other constraints on prices of tranches?
- Should the price of a combined tranche be equal to the sum of prices of all single tranches, which constitute a combined tranche?
- Under which assumptions, a consistent valuation of all tranches (hence of a multi-person game claim) is feasible?

4.1 Super-Additivity of Lower Prices

We first show that the lower prices $\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ are inherently super-additive.

Proposition 4.1. *The lower price of* $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$*, which is defined by*

$$\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) = \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma),$$

has the following super-additive property: if $\mathcal{N}_1, \mathcal{N}_2, \ldots, \mathcal{N}_k$ *is a partition of* $\mathcal{N}$*, then*

$$\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \geq \sum_{i=1}^{k} \underline{\pi}_t^{\mathcal{N}_i}(\mathcal{G}, s).$$

Proof. According to Lemma 3.11, for each of the sets $\mathcal{N}_i$, we may choose a sequence $(\tau_n^{\mathcal{N}_i})_{n \in \mathbb{N}}$ from $\mathcal{S}_t^{\mathcal{N}_i}(s)$ such that the sequence

$$\operatorname*{ess\,inf}_{\sigma^{-\mathcal{N}_i} \in \mathcal{S}_t^{-\mathcal{N}_i}(s)} \widehat{V}_t^{\mathcal{N}_i}\left(\tau_n^{\mathcal{N}_i}, \sigma^{-\mathcal{N}_i}\right), \quad n \in \mathbb{N}$$

is non-decreasing and converges to $\underline{\pi}_t^{\mathcal{N}_i}(\mathcal{G}, s)$. Now define the strategy $\tau_n \in \mathcal{S}_t^{\mathcal{N}}(s)$ for $n \in \mathbb{N}$ by $\tau_n := (\tau_n^{\mathcal{N}_1}, \ldots, \tau_n^{\mathcal{N}_k})$. For each $i = 1, \ldots, k$ and for all $\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)$, we have

$$\widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma) = \sum_{i=1}^{k} \widehat{V}_t^{\mathcal{N}_i}(\tau_n, \sigma) \geq \sum_{i=1}^{k} \operatorname*{ess\,inf}_{\sigma^{-\mathcal{N}_i} \in \mathcal{S}_t^{-\mathcal{N}_i}(s)} \widehat{V}_t^{\mathcal{N}_i}\left(\tau_n^{\mathcal{N}_i}, \sigma^{-\mathcal{N}_i}\right).$$

Since this holds for all $\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)$, the above inequality implies that

$$\operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma) \geq \sum_{i=1}^{k} \operatorname*{ess\,inf}_{\sigma^{-\mathcal{N}_i} \in \mathcal{S}_t^{-\mathcal{N}_i}(s)} \widehat{V}_t^{\mathcal{N}_i}\left(\tau_n^{\mathcal{N}_i}, \sigma^{-\mathcal{N}_i}\right)$$

and thus

$$\begin{aligned}\underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) &= \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma) \geq \sup_{n \in \mathbb{N}} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau_n, \sigma) \\ &\geq \sup_{n \in \mathbb{N}} \sum_{i=1}^{k} \operatorname*{ess\,inf}_{\sigma^{-\mathcal{N}_i} \in \mathcal{S}_t^{-\mathcal{N}_i}(s)} \widehat{V}_t^{\mathcal{N}_i}\left(\tau_n^{\mathcal{N}_i}, \sigma^{-\mathcal{N}_i}\right) = \sum_{i=1}^{k} \underline{\pi}_t^{\mathcal{N}_i}(\mathcal{G}, s),\end{aligned}$$

which completes the proof. □

The super-additive property from Proposition 4.1 can be easily explained as follows. First note that, by Lemma 3.11, the holder of any individual tranche may choose a strategy to guarantee an expected discounted payoff which is arbitrarily close to its lower price. So by playing these strategies over a collection of tranches, one can guarantee a total expected discounted payoff arbitrarily close to the sum of individual lower prices, which is in turn no greater than the lower price of the collection of tranches.

Unfortunately, no analogous additive properties exist for the upper prices $\overline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$. Furthermore, from the analysis so far, there is no reason for the prices $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$ themselves to satisfy any additive properties either. Does that mean the bounds described by (3.34) are the only important constraints when pricing game contingent claims?

4.2 Additivity of Prices

As illustrated by Example 3.1, even if the bounds of (3.34) hold (and thus the no-arbitrage property of Definition 3.1 is satisfied), they do not necessarily lead to sensible pricing of the contract. There are two main causes of this discrepancy. First, each party simply cannot rule out the possibilities of the others colluding. This perceived possibility of collusion can also be interpreted as the uncertainty over the identity of the other holders, that is, multiple holders (or issuer) may actually be the same person. Second, the arbitrage prices in Example 3.1 do not satisfy a natural requirement of price additivity.

Definition 4.1. We say that the *price additivity* holds whenever

$$\pi_t^{\mathcal{N}}(\mathcal{G}, s) = \sum_{k \in \mathcal{N}} \pi_t^k(\mathcal{G}, s), \quad \forall \mathcal{N} \subseteq \mathcal{M}.$$

In general, the problem highlighted by Example 3.1 occurs under the following conditions. Suppose there are k prospective buyers, who intend to purchase tranches $\mathcal{N}_1, \ldots, \mathcal{N}_k$, respectively, while the issuer decides to keep the tranche $-\mathcal{N} = -(\mathcal{N}_1 \cup \cdots \cup \mathcal{N}_k)$. It is not possible to reach agreements on prices if the lower price of the tranche $\mathcal{N}$ exceeds the sum of upper prices of the tranche $\mathcal{N}_i$,

$$\sum_{i=1}^{k} \overline{\pi}_t^{\mathcal{N}_i}(\mathcal{G}, s) < \underline{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) \quad \Longrightarrow \quad \sum_{i=1}^{k} \pi_t^{\mathcal{N}_i}(\mathcal{G}, s) < \pi_t^{\mathcal{N}}(\mathcal{G}, s). \tag{4.1}$$

Indeed, the issuer will then always demand more than the what the holders are willing to pay. However, if the price additivity holds, then strict inequality (4.1) cannot occur.

Example 3.1 offered some justification to why the additivity of prices is a desirable condition. To further consolidate this belief, we shall introduce a second form of arbitrage called *immediate reselling arbitrage*. In the previous section, Theorems 3.1 and 3.3 prevent enter-and-hold arbitrage opportunities (as per Definition 3.1), which are executed by holding a fixed collection of tranches from time t until the settlement time. An immediate reselling arbitrage simply involves buying some tranches at time t and then instantly reselling them in a different configuration with the goal of collecting a higher total price. We will define this formally.

Definition 4.2. Consider the game contingent claim $\mathrm{MGC}_t(\mathcal{G}, s)$ at time t. For each $\mathcal{N} \subseteq \mathcal{M}$, denote the price of the tranche $\mathcal{N}$ by $\pi_t^{\mathcal{N}}(\mathcal{G}, s)$. A *immediate reselling arbitrage* consists of two different partitions of some subset $\mathcal{N} \subseteq \mathcal{M}$,

$$\mathcal{N} = \mathcal{N}_1 \cup \cdots \cup \mathcal{N}_k = \mathcal{N}_1' \cup \cdots \cup \mathcal{N}_l',$$

such that the inequality

$$\sum_{i=1}^{k} \pi_t^{\mathcal{N}_i}(\mathcal{G}, s) < \sum_{j=1}^{l} \pi_t^{\mathcal{N}_j'}(\mathcal{G}, s)$$

holds on some $\mathcal{F}_t$-measurable event A with a positive probability.

In European and American option pricing models, immediate reselling arbitrages are implicitly avoided due to the law of one price. In multi-person game contingent claims, the ability to purchase collections of tranches has reintroduced this possibility. However, it is easy to see that immediate reselling arbitrages are avoided whenever the price additivity holds. The following result is an immediate consequence of Definition 4.2.

Proposition 4.2. *There is no immediate reselling arbitrage if and only if the additivity of prices holds. In other words,*

$$\pi_t^{\mathcal{N}}(\mathcal{G}, s) = \sum_{k \in \mathcal{N}} \pi_t^k(\mathcal{G}, s), \quad \forall \mathcal{N} \subseteq \mathcal{M}.$$

By combining Proposition 4.2 with Theorems 3.1 and 3.3, we obtain the following characterisation of prices which avoids both types of arbitrage mentioned. As usual, we consider a multi-person stochastic game whose payoffs satisfy the appropriate conditions (Assumption 3.1 for the discrete-time set-up and Assumption 3.1 for the continuous-time set-up).

Theorem 4.1. *For every $\mathcal{N} \subseteq \mathcal{M}$, let $\widehat{\pi}_t^{\mathcal{N}}(\mathcal{G}, s)$ be the discounted time t price of the combined tranche $\mathcal{N}$ in $\mathrm{MGC}_t^{\mathcal{N}}(\mathcal{G}, s)$. Then there is no arbitrage in $(B, S, \mathcal{N})$ for all $\mathcal{N} \subseteq \mathcal{M}$ and no immediate reselling arbitrage if and only if the following two conditions are satisfied:*
(i) the equality

$$\widehat{\pi}_t^{\mathcal{N}}(\mathcal{G}, s) = \sum_{k \in \mathcal{N}} \widehat{\pi}_t^k(\mathcal{G}, s)$$

holds for all $\mathcal{N} \subseteq \mathcal{M}$,
(ii) the random vector $(\widehat{\pi}_t^1(\mathcal{G}, s), \ldots, \widehat{\pi}_t^m(\mathcal{G}, s))$ lies in the following random subset of $\mathbb{R}^m$,

$$\Big\{(X_1, \ldots, X_m) : \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma) \leq \sum_{i \in \mathcal{N}} X_i \tag{4.2}$$
$$\leq \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-\mathcal{N}}(s)} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^{\mathcal{N}}(s)} \widehat{V}_t^{\mathcal{N}}(\tau, \sigma), \quad \forall \mathcal{N} \subseteq \mathcal{M}\Big\},$$

where $X_1, \ldots, X_m$ are $\mathcal{F}_t$-measurable random variables.

Remark 4.1. In some game contingent claims $\mathrm{MGC}(\mathcal{G})$, such as Example 3.1, the conditions of Theorem 4.1 cannot be satisfied because the set in (4.2) is empty. This means that some game contingent claims $\mathrm{MGC}(\mathcal{G})$ simply do not admit prices under our arbitrage-free assumptions. Occasionally this may be remedied if the identities of the holders (or the possibility of collusion) is known to everyone in advance, but in these cases $\mathrm{MGC}(\mathcal{G})$ is equivalent to a (simpler) contract with less tranches. Another solution is to alter the definitions of arbitrage and weaken some of the pricing bounds. In any case, it is fair to admit that condition (4.2) is not handy and thus we will present more explicit conditions in Theorem 4.2 below.

4.3 Optimal Equilibrium and Arbitrage Prices

From a purely game-theoretic perspective, the upper and lower prices are also the minimax and maximin values of the coalition $\mathcal{N}$ in the subgame of $\mathcal{G}$ on $[t, T]$. Note that the expected discounted payoff under the martingale measure $\mathbb{Q}$ is used in this case. As shown in Section 3.7, the properties of minimax and maximin strategies may be interpreted using super-hedging arguments.

In general, the infimum in the upper price and the supremum in the lower price are not necessarily achieved. But if $\mathcal{G}$ is known to have an optimal equilibrium s^* then, from Corollary 5.1, the maximin and minimax values for each player (or equivalently, the lower and upper prices of individual tranches) are achieved by $\widehat{V}_t(s^*)$, i.e., by the unique value of the game. Consequently, the vector of prices of individual tranches must coincide with the unique value of the game.

As demonstrated by means of Example 3.1, the mere existence of optimal equilibria still does not lead to additive prices. However, Proposition 5.3 in the appendix shows that if the game $\mathcal{G}$ is

also known to satisfy some further conditions (e.g., zero-sum), then coalition values exist and satisfy the additivity property. In that case, the prices of combined tranches are also unique and the price additivity holds.

We summarise these results in the following theorem. Note that the date $t \in [0,T)$ is fixed, but arbitrary. We work under our standing assumptions, that is, Assumption 3.1 for the discrete-time case and Assumption 3.1 for the continuous-time case.

Theorem 4.2. *Consider a game $\mathcal{G}$ and the associated game contingent claim* $\mathrm{MGC}(\mathcal{G})$. *Assume that the strategy profile $s_{[0,t)} \in \mathcal{S}_{[0,t)}$ is chosen before time t and $\vartheta(s) \geq t$. Let the expected discounted payoff under the unique martingale measure $\mathbb{Q}$ be denoted by $\widehat{V}_t(\tau,\sigma) = \mathbb{E}_{\mathbb{Q}}\big(\widehat{V}_\vartheta(\tau,\sigma) \,\big|\, \mathcal{F}_t\big)$.*

Suppose that the subgame on the time interval $[t,T]$ has an optimal equilibrium $s^* = (\tau^*,\sigma^*) \in \mathcal{S}_t(s)$, meaning that

$$\widehat{V}^k_t(\tau^*,\sigma^*) = \operatorname*{ess\,inf}_{\sigma \in \mathcal{S}_t^{-k}(s)} \widehat{V}^k_t(\tau^*,\sigma) = \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t^k(s)} \widehat{V}^k_t(\tau,\sigma^*), \quad \forall\, k \in \mathcal{M}. \tag{4.3}$$

(i) Then, for any $k \in \mathcal{M}$, there is no arbitrage in (B,S,k) if and only if the discounted time t price of a single tranche k satisfies

$$\widehat{\pi}^k_t(\mathcal{G},s) = \widehat{V}^k_t(s^*).$$

(ii) Assume, in addition, that the optimal equilibrium s^* satisfies the following condition

$$\sum_{k\in\mathcal{M}} \widehat{V}^k_t(s^*) = \operatorname*{ess\,sup}_{s' \in \mathcal{S}_t(s)} \sum_{k\in\mathcal{M}} \widehat{V}^k_t(s'). \tag{4.4}$$

Then, for any $\mathcal{N} \in \mathcal{M}$, there is no arbitrage in $(B,S,\mathcal{N})$ if and only if the discounted time t price of a combined tranche $\mathcal{N}$ has the following additivity property

$$\widehat{\pi}^{\mathcal{N}}_t(\mathcal{G},s) = \sum_{k\in\mathcal{N}} \widehat{V}^k_t(s^*) = \widehat{V}^{\mathcal{N}}_t(s^*). \tag{4.5}$$

Furthermore, there is no immediate reselling arbitrage if and only if (4.5) *is satisfied for all $\mathcal{N} \subseteq \mathcal{M}$.*

Proof. (i) Corollary 5.1 from the appendix shows that all optimal equilibria achieve the minimax and maximin values simultaneously. Since we know from Theorems 3.1 and 3.3 that an arbitrage price must lie between the upper and lower prices, which are identical to the minimax and maximin values, an arbitrage price must in fact coincide with the unique value achieved by any optimal equilibrium.

(ii) Proposition 5.3 from the appendix states that, under (4.4), the coalition values exist and enjoy the additivity property. Formula (4.5) thus follows immediately from Theorems 3.1 and 3.3, since the arbitrage price of a combined tranche must coincide with the coalition value. The lack of the immediate reselling arbitrage is a consequence of Proposition 4.2 and the price additivity. □

4.4 Classes of Solvable Multi-Person Games

Part (ii) in Theorem 4.2 is of great significance, since it further justifies the choice of the optimal equilibrium as a solution concept when studying multi-person game contingent claims, as it produces unique arbitrage prices for individual tranches. In general, as shown in Section 3.6, this would not be possible to achieve with weaker solution concepts, such as the Nash equilibrium.

Let us thus focus on multi-player games that admit an optimal equilibrium under the unique martingale measure $\mathbb{Q}$. In this situation, the existence of unique and additive arbitrage prices for single tranches of a multi-person game claim hinges on condition (4.4). Hence the following natural problem arises: describe sufficiently large classes of multi-player games, which admit an optimal equilibrium and satisfy condition (4.4).

This issue is of interest even in the case of two-person non-zero-sum Dynkin games, although it is fair to say that in most existing papers devoted to Dynkin games, the authors focus on the existence of a Nash equilibrium (see, for instance, Hamadène and Hassani [14, 15], Hamadène and Zhang [16], Laraki and Solan [28], Peskir [32] or Ohtsubo [30]). In particular, in [16], the authors examine indifference pricing of financial contracts, which are based on non-zero-sum Dynkin games. This should be contrasted with the case of a zero-sum Dynkin game, for which the arbitrage valuation theory can be applied in a relatively straightforward manner, as demonstrated by Kifer [24] (for related results in this vein, see also [6, 7, 20, 21, 25, 26, 27]).

For the multi-player case with $m \geq 3$, an explicit specification of a sensible game admitting an optimal equilibrium is a challenging issue, which was studied in some recent works (see [12, 13, 29]). Due to limited space, we do not discuss this issue in detail here, and we instead restrict ourselves to providing some examples of multi-person games with an optimal equilibrium satisfying condition (4.4). In particular, in Section 4.5, we solve a particular valuation problem for a multi-person contract with puttable tranches, which was described in Section 1.1.

As a first example of a game admitting an optimal equilibrium, we may quote any two-player zero-sum game in which only one player has meaningful decisions to make or, equivalently, any single player game. In fact, this class of relatively simple stochastic games covers most of conventional financial derivatives in which the 'issuer' has no right to make any decisions about a sold contract. Obviously, we have in mind here plain-vanilla American options, for which the valuation hinges on solving an optimal stopping problems, but also more complex derivatives of an American style with some 'exotic' features, which may allow the holder to make other relevant decisions, not only exercising (e.g., an American chooser option). Original examples of more sophisticated multi-player games satisfying the assumptions of part (ii) in Theorem 4.2 are collected in the following proposition.

Proposition 4.3. *Assume that a game $\mathcal{G}$ belongs to either of the following classes of multi-player games:*
(i) zero-sum (or constant-sum) games that admits an optimal equilibrium,
(ii) multi-period redistribution games, denoted as $\mathrm{MRG}(X, \alpha)$*,*
(iii) affine stopping games, denoted as $\mathrm{ASG}(X, G)$*, with a $\mathcal{K}^\dagger$-matrix G with non-negative column sums.*
Then $\mathcal{G}$ has an optimal equilibrium satisfying condition (4.4)

Proof. (i) This part is rather obvious, since both sides in (4.4) vanish identically (or are equal to the same constant) for a zero-sum (or a constant-sum) game.

(ii) Multi-period redistribution games $\mathrm{MRG}(X, \alpha)$ are studied in [12], where it was shown that these zero-sum (non-zero-sum, resp.) redistribution games have an optimal equilibrium by virtue of Theorem 2 (Theorem 3, resp.).

(iii) Recall that $\mathcal{K}^\dagger$-matrix is a square matrix with a non-negative determinant, positive proper principal minors, and with non-positive off-diagonal terms (see Definition 3.2 in [13]). The desired properties of an affine stopping game $\mathrm{ASG}(X, G)$ with a $\mathcal{K}^\dagger$-matrix G were established in Theorem 3.2 in [13]. □

In the same way a general game contingent claim is defined, one can define a multi-person game contingent claim based on a multi-player game $\mathcal{G}$. Of course, a classic example of this approach is the definition of the two-person game claim, which corresponds to a zero-sum two-person Dynkin game. It should be mentioned, however, that in a multi-period (or continuous-time) setup, the redistribution and affine games introduced in Guo [10] and Guo and Rutkowski [11, 12, 13] are defined recursively, meaning that the payoffs at settlement depend on the value of analogous virtual game continued after the settlement date (this concept is reminiscent of a continuation value for an American or game claim). For this reason, the optimal equilibria are obtained either through the backward induction or by solving a multi-dimensional reflected BSDE (see Nie and Rutkowski [29]). The latter approach hinges on a multi-dimensional extension of the one-dimensional doubly-reflected BSDE for the Dynkin game studied in Cvitanić and Karatzas [5].

4.5 Pricing of Multi-Person Contracts with Puttable Tranches

We conclude this work by examining arbitrage pricing of a multi-person contract with puttable tranches. Recall that the mechanism of this contract was outlined in Section 1.1. The key idea is to reduce the valuation problem into the search of optimal equilibria for the associated multi-player stochastic game under a unique martingale measure for the underlying market model. We then apply the results from our previous papers [11, 12] to find the solution to the game. In fact, the results from [12] need to be slightly extended, since the redistribution games studied in [12] are formally postulated to terminate as soon as at least one of the players decides to exercise (i.e., stop the game), whereas in Section 1.1 we introduced a contract, which always continues till its maturity date $T = T_{n+1}$. Note that we work here with either a discrete- or continuous-time setting.

Let us consider the date T_n. Suppose that there exists a unique martingale measure $\mathbb{Q}$ for the underlying market model, say, the CRR binomial model or the Black and Scholes model. We assume, for simplicity, that the interest rate is null (otherwise, one would work with the discounted values). Then the prices $P_n^1, \ldots, P_n^m$ of tranches at time T_n, before decisions of all parties at time T_n were made (dubbed the *continuation values* at time T_n), are prices of European or American options and thus they can be found by the usual arbitrage pricing method independently for each tranche. For instance, if the terminal payoffs of European contingent claims are $X_{n+1}^1, \ldots, X_{n+1}^m$, then we obtain the equality $P_n^i = \mathbb{E}_{\mathbb{Q}}(X_{n+1}^i \,|\, \mathcal{F}_{T_n})$ for every i where $\mathbb{F} = (\mathcal{F}_t)_{t\in[0,T]}$ is the underlying filtration.

A strategy for the holder of tranche i on $[0, T_n]$ is denoted by s^i. In the present setting, it consists of choosing between 'hold' or 'put' at each date $T_1, \ldots, T_n$. These decisions are made by the current holder of tranche i as time progresses, which may not be the same person, due to the nature of the contract. Let us denote a generic strategy profile by s and its restriction to $[T_l, T_n]$ by s_l. It will be convenient to represent a strategy profile s_l as $s_l = (\tau_l, \sigma_l)$ where $\tau_l = s_l^i$ ($\sigma_l = s_l^{-i}$, resp.) is the strategy of player i (the strategy profile for all other players, resp.) on $[T_l, T_n]$. We denote by $\mathcal{E}_l(s_l)$ the set of players who decide to put the tranche at time T_l when a strategy profile s_l is played on $[T_l, T_n]$. To proceed further, we need to introduce some notation, which is adapted from [12]. For any $X = (X^1, \ldots, X^m) \in \mathbb{R}^m$ and an arbitrary subset $\mathcal{E} \subseteq \mathcal{M}$, we define the hyperplane

$$\mathbb{H}(\mathcal{E}, X, c) := \left\{ x \in \mathbb{R}^m : \sum_{i=1}^m x_i = c \text{ and } x_i = X^i,\ \forall\, i \in \mathcal{E} \right\}$$

and the simplex

$$\mathbb{S}(X, c) := \left\{ x \in \mathbb{R}^m : \sum_{i=1}^m x_i = c \text{ and } x_i \geq X^i,\, \forall\, i \in \mathcal{M} \right\}.$$

Let an $\mathcal{F}_{T_l}$-measurable random variable $(C_l^1, \ldots, C_l^m)$ and an $\mathcal{F}_{T_l}$-measurable subset $\mathcal{E}_l$ of $\mathcal{M}$ be given. Then we define the random, $\mathcal{F}_{T_l}$-measurable hyperplane $\mathbb{H}(\mathcal{E}_l, X_l, c_l)$ and the random $\mathcal{F}_{T_l}$-measurable simplex $\mathbb{S}(X_l, c_l)$, where $c_l := \sum_{i=1}^m C_l^i$. Let us finally observe that, since for the uniform redistribution we have that $\alpha_i = \frac{1}{m}$ for all i, the inner product given by equation (3.9) in [12] reduces to the Euclidean inner product and thus the projection π^0 utilised in Theorem 3.12 in [12] is the usual orthogonal projection in $\mathbb{R}^m$; it will be henceforth denoted as π. The next definition hinges on Lemma 3.11 in [12], which furnishes an explicit characterisation (4.6) of the effective (or *modified*) payoff in a zero-sum single-period redistribution game. We denote by V_{t+1}^* the *value* of the subgame started at T_{t+1} and thus the definition of the game is recursive (its correctness is formally supported by Proposition 4.4 where the existence and uniqueness of the value process is established).

Definition 4.3. Assume that an $\mathbb{F}$-adapted process $X_l = (X_l^1, \ldots, X_l^m)$, $l = 1, \ldots, n$ and the continuation value $P_n = (P_n^1, \ldots, P_n^m)$ at time T_n are given. Let $V_l^i(s_l)$ denote the $\mathcal{F}_{T_l}$-conditional expected risk-neutral payoff at time T_l for the current holder of tranche i given that a strategy profile s_l is played on $[T_l, T_n]$. In the *multi-player game with puttable tranches*, for any strategy profile s, $V_l(s_l)$ is defined recursively through the following expression, for every $l = 1, \ldots, n$ and $i = 1, \ldots, m$,

$$V_l^i(s_l) := X_l^i \mathbb{1}_{\{\vartheta^i(s_l)=l\}} + \left(\left[\pi_{\mathbb{H}(\mathcal{E}_l(s_l), X_l, p_l)}(P_l)\right]^i - P_l^i + \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^i(s_{l+1}) \,|\, \mathcal{F}_{T_l}\big) \right) \mathbb{1}_{\{\vartheta^i(s_l)>l\}} \tag{4.6}$$

where the equality $\vartheta^i(s_l) = l$ (the inequality $\vartheta^i(s_l) > l$, resp.) means that the holder of the ith tranche decides to put (decides not to put, resp.) at time T_l and

$$p_l := \sum_{i=1}^{m} P_l^i \tag{4.7}$$

where $P_l^i := \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^{i*} \,|\, \mathcal{F}_{T_l}\big)$ for every $i = 1, \ldots, m$ and $l = 1, \ldots, n-1$. Furthermore, we set $V_0(s) = \mathbb{E}_{\mathbb{Q}}(V_1(s_1))$ since putting a tranche at time 0 is not allowed.

Observe that the *continuation value* of tranche i at time T_l is given as $P_l^i = \mathbb{E}_{\mathbb{Q}}(V_{l+1}^{i*} \,|\, \mathcal{F}_{T_l})$ for every i. In equation (4.6), the term X_l^i represents the payoff received by holder of tranche i when he decides to put at time T_l, and the quantity in front of the indicator $\mathbb{1}_{\{\vartheta^i > l\}}$ represents the expected payoff of holder when he decides not to put the tranche. Upon summation, we can represent the right-hand side in (4.6) as follows

$$V_l^i(s_l) := \mathbb{E}_{\mathbb{Q}}\bigg(X_{\vartheta^i(s_l)}^i + \sum_{t=l}^{\vartheta^i(s_l)-1} \Big[\pi_{\mathbb{H}(\mathcal{E}_t(s_t), X_t, p_t)}\big(\mathbb{E}_{\mathbb{Q}}\big(V_{t+1}^* \,|\, \mathcal{F}_{T_t}\big)\big) - \mathbb{E}_{\mathbb{Q}}\big(V_{t+1}^* \,|\, \mathcal{F}_{T_t}\big)\Big]^i \,\bigg|\, \mathcal{F}_{T_l}\bigg) \tag{4.8}$$

where $\vartheta^i(s_l)$ represents the index such that the random time $T_{\vartheta^i(s_l)} \geq T_l$ is the first instance at which tranche i is put. Hence the summation term in (4.8) represents the accumulated redistribution holder i is required to pay as a result of other holders putting their tranches during $[T_l, T_{\vartheta^i(s_l)})$.

It should be noted that game described in Definition 4.3 is nothing more than a suitable variant of the game given by Definition 4.1 in [12]. The only difference is that, in [12], the multi-person game is stopped as soon as at least one player decides to put his tranche. In that case, the expected payoff is determined using, in particular, the value process of a virtual game played from this moment onwards. Therefore, Definition 4.1 in [12] refers in fact to a family of the embedded games that start at any date T_l; each of them ends as soon as someone 'exercises' (i.e., a holder puts his tranche, using the present terminology). By contrast, in Definition 4.3, we formally specify a single game, which is always played till expiration date T, but is otherwise specified in the same way as the game studied in [12]. It is thus natural to conjecture that the two games share the same value process and an optimal equilibrium. The validity of this conjecture will be established in Proposition 4.4. The proof of this result utilises Theorem 1 in [11] (see also Theorem 3.12 in [12]).

We acknowledge that the game introduced in Definition 4.3 departs from the usual way in which multi-person extensions of classic Dynkin games are defined (see, for instance, Hamadène and Hassani [14, 15] or Karatzas and Li [22]), since we postulate that the expected payoff at time T_l depends not only on exogenously specified payoff processes $X^1, \ldots, X^m$, but also on the value of the embedded game started at T_l. This way of defining a game could seem unusual in the context of the game theory, but it is perfectly consistent with the arbitrage pricing of derivative securities and the concept of the market value, where the future decisions of investors need to be taken into account. Also, in the case of a two-person game option, the corresponding Dynkin game is formally defined using only the payoff processes, say X^1 and X^2. However, an equivalent formulation that utilises processes $X^1 \geq X^2$ and the continuation value P of the contract (or, equivalently, the associated Dynkin game) is also possible and in fact it provides a better insight into competitive character of a game option. However, since the gain of one party matches the loss of the counterparty, the trivial mechanism of redistribution is already implicit in the two-person version of the game. Of course, when $m = 2$, the game of Definition 4.3 reduces to the classic Dynkin game and a multi-person contract with puttable tranches becomes a two-person game option.

Since it is possible for the issuer to simultaneously hold multiple tranches at one point, the price additivity (see Definition 4.1) becomes an important issue. We hence need to check a sub-zero-sum condition on the value V_0^*. The most convenient solution is to enforce a restriction such as (4.9), so that it will never be optimal for all holders to simultaneously put, and thus the contract will be a zero-sum game between the holders.

To ensure this property in the case of a contract described in Section 1.1, one should choose strikes $K_1, \ldots, K_m$, classes of options (calls or puts), and the *put payoffs* $X_l^1, \ldots, X_l^m$, $l = 1, \ldots, n$ in such a way that it will be not 'optimal' (in the sense of maximisation of the value of a tranche) for all holders to simultaneously give their tranches back to the issuer at some date T_l.

Proposition 4.4. *Assume that the zero-sum condition*

$$\sum_{i=1}^{m} X_l^i \leq \sum_{i=1}^{m} P_l^i \tag{4.9}$$

is satisfied recursively for $l = 1, \ldots, n$. Then the multi-player game with puttable tranches admits an optimal equilibrium s^* and a unique value process $V^* = V(s^*)$.

Proof. In the first step, we consider the penultimate date T_n. Then the continuation values $P_n^1, \ldots, P_n^m$ are obtained by computing the continuation values for all tranches at time T_n using the standard arbitrage pricing method. Consequently, the multi-person game with puttable tranches can be reduced at time T_n to a single-period zero-sum redistribution game $\mathrm{ZRG}(X_n, P_n, \alpha)$ with $\alpha_i = \frac{1}{m}$ for all i (see Definitions 3.1 and 3.10 in [12]). Therefore, the existence of an optimal equilibrium $s_n^* = (\tau_n^*, \sigma_n^*)$ (hence also of a unique value V_n^* for the game at time T_n) follows from Theorem 1 in [11] (see also Theorem 3.12 in [12]). Specifically, the unique value for each individual tranche at time T_n equals

$$V_n^* = V_1(s_n^*) = \pi_{\mathbb{S}(X_n, p_n)}(P_n)$$

provided that

$$\sum_{i=1}^{m} X_n^i \leq p_n = \sum_{i=1}^{m} P_n^i,$$

that is, it is not optimal for all players to put tranches at time T_n.

We now proceed by backward induction. We fix $l = 1, \ldots, n-1$, and we suppose that the value at time T_{l+1} is $V_{l+1}^* = V_{l+1}(s_{l+1}^*)$ where a strategy profile s_{l+1}^* is an optimal equilibrium for the game on the time interval $[T_{l+1}, T_n]$, meaning the following inequalities hold for every i and all strategies τ_{l+1} and σ_{l+1}

$$V_{l+1}^i(\tau_{l+1}^*, \sigma_{l+1}) \geq V_{l+1}^{*i} \geq V_{l+1}^i(\tau_{l+1}, \sigma_{l+1}^*). \tag{4.10}$$

Recall that the vector of continuation values at time T_l is given by $P_l^i := \mathbb{E}_{\mathbb{Q}}(V_{l+1}^{*i} \,|\, \mathcal{F}_{T_l})$ for all i. Let us then suppose that the value V_{l+1}^* of the game at time T_{l+1}, and thus also the continuation values $P_l^1, \ldots, P_l^m$ of tranches, before players' decisions at time T_l were made, are given.

We define the strategy profile s_l^* on $[T_l, T_n]$ by postulating that: (i) the set of players who decide to put the tranche at the time T_l equals

$$\mathcal{E}_l(s_l^*) = \big\{ i \in \{1, \ldots, m\} : \big[\pi_{\mathbb{S}(X_l, p_l)}(P_l)\big]^i = X_l^i \big\} \tag{4.11}$$

and (ii) s_l^* restricted to $[T_{l+1}, T_n]$ coincides with s_{l+1}^*. Using (4.8)–(4.7) and property (ii), we obtain

$$V_l(s_l^*) = \pi_{\mathbb{H}(\mathcal{E}_l(s_l^*), X_l, p_l)}(P_l) - P_l + \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}(s_{l+1}^*) \,|\, \mathcal{F}_{T_l}\big) = \pi_{\mathbb{H}(\mathcal{E}_l(s_l^*), X_l, p_l)}(P_l)$$

where the set $\mathcal{E}_l(s_l^*)$ is given by (4.11) and p_l is given by (4.7). We aim to demonstrate that the strategy profile s_l^* is an optimal equilibrium on $[T_l, T_n]$ and $V_l^* := V_l(s_l^*)$ is the game value at time T_l. To this end, we need to show that, for every i,

$$V_l^i(\tau_l^*, \sigma_l) \geq V_l^i(s_l^*) = V_l^i(\tau_l^*, \sigma_l^*) \geq V_l^i(\tau_l, \sigma_l^*) \tag{4.12}$$

for arbitrage strategy profiles τ_l and σ_l on $[T_l, T_n]$, not necessarily coinciding with τ_{l+1}^* and σ_{l+1}^* on $[T_{l+1}, T_n]$.

From Theorem 1 in [11] (or Theorem 3.12 in [12]), it is known that condition (i) characterises an optimal equilibrium in the single-period game $\mathrm{ZRG}(X_l, P_l, \alpha)$ and thus

$$V_l(s_l^*) = \pi_{\mathbb{H}(\mathcal{E}_l(s_l^*),X_l,p_l)}(P_l) = \pi_{\mathbb{S}(X_l,p_l)}(P_l) \tag{4.13}$$

provided that $\sum_{i=1}^m X_l^i \leq p_l$. To be more specific, for every $i = 1, \ldots, m$ and all strategies τ_l and σ_l such that $\tau_{l+1} = \tau_{l+1}^*$ and $\sigma_{l+1} = \sigma_{l+1}^*$, we have

$$\begin{aligned} \bar{V}_l^i(\tau_l^*, \sigma_l) &= \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l^*,\sigma_l),X_l,p_l)}(P_l)\big]^i \geq [\pi_{\mathbb{S}(X_l,p_l)}(P_l)]^i \\ &\geq \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l,\sigma_l^*),X_l,p_l)}(P_l)\big]^i = \bar{V}_l^i(\tau_l, \sigma_l^*) \end{aligned} \tag{4.14}$$

where we write $\bar{V}_l^i$, rather than V_l^i, to emphasise that we deal here with a single-period game on $[T_l, T_{l+1}]$, which is obtained by postulating that all strategy profiles that appear in the formula above necessarily coincide with s_{l+1}^* on $[T_{l+1}, T_n]$.

To show the Nash equilibrium property of s_l^* at time T_l for the multi-player game given by Definition 4.3, let us take any strategy profile τ_l for the holder of tranche i. If τ_l implies putting tranche i at time T_l, then from (4.14), we obtain

$$V_l^i(\tau_l, \sigma_l^*) = X_l^i = \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l,\sigma_l^*),X_l,p_l)}(P_l)\big]^i \leq [\pi_{\mathbb{S}(X_l,p_l)}(P_l)]^i = V_l^i(s_l^*).$$

If τ_l does not imply putting tranche i at time T_l, then from (4.10), (4.13) and (4.14), we obtain

$$\begin{aligned} V_l^i(\tau_l, \sigma_l^*) &= \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l,\sigma_l^*),X_l,p_l)}(P_l)\big]^i - P_l^i + \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^i(\tau_{l+1}, \sigma_{l+1}^*) \,|\, \mathcal{F}_{T_l}\big) \\ &\leq [\pi_{\mathbb{S}(X_l,p_l)}(P_l)]^i - P_l^i + \mathbb{E}_{\mathbb{Q}}(V_{l+1}^{*i} \,|\, \mathcal{F}_{T_l}) = V_l^i(s_l^*). \end{aligned}$$

To verify the guaranteed payoff inequality, we take an arbitrary strategy profile σ_l and we use similar arguments as above. If τ_l^* implies putting tranche i at time T_l, then from (4.14), we obtain

$$V_l^i(\tau_l^*, \sigma_l) = X_l^i = \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l^*,\sigma_l),X_l,p_l)}(P_l)\big]^i \geq [\pi_{\mathbb{S}(X_l,p_l)}(P_l)]^i = V_l^i(s_l^*).$$

If τ_l^* does not imply putting tranche i at time T_l, then from (4.10), (4.13) and (4.14), we obtain

$$\begin{aligned} V_l^i(\tau_l^*, \sigma_l) &= \big[\pi_{\mathbb{H}(\mathcal{E}_l(\tau_l^*,\sigma_l),X_l,p_l)}(P_l)\big]^i - P_l^i + \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^i(\tau_{l+1}^*, \sigma_{l+1}) \,|\, \mathcal{F}_{T_l}\big) \\ &\geq [\pi_{\mathbb{S}(X_l,p_l)}(P_l)]^i - P_l^i + \mathbb{E}_{\mathbb{Q}}(V_{l+1}^{*i} \,|\, \mathcal{F}_{T_l}) = V_l^i(s_l^*). \end{aligned}$$

We have thus shown that (4.12) holds for every i and all strategies τ_l and σ_l. This means that s_l^* is indeed an optimal equilibrium on $[T_l, T_n]$ and V_l^* is the value of the game at time T_l. Finally, the value of the game at time 0 equals $V_0^* = \mathbb{E}_{\mathbb{Q}}(V_1^*)$, since putting a tranche at time 0 is not allowed. □

One could also consider a simplified version of the game of Definition 4.3, which is specified as follows: we assume that the vector $P_n = (P_n^1, \ldots, P_n^m)$ is given and, for any strategy profile s on $[0, T_n]$, we set

$$V_n(s_n) := \pi_{\mathbb{H}(\mathcal{E}_n(s_n),X_n,p_n)}(P_n) \tag{4.15}$$

and we define $V_l(s_l)$ recursively, for every $l = 1, \ldots, n-1$,

$$V_l(s_l) := \pi_{\mathbb{H}(\mathcal{E}_l(s_l),X_l,p_l)}\big(\mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^* \,|\, \mathcal{F}_{T_l}\big)\big)$$

where V_{l+1}^* is the value of the subgame started at T_{l+1} and p_l, $l = 1, \ldots, n-1$ is given by (4.7) with $P_l^i := \mathbb{E}_{\mathbb{Q}}\big(V_{l+1}^{*i} \,|\, \mathcal{F}_{T_l}\big)$. For this game, Proposition 4.4 is still valid and its proof can be essentially simplified.

It is less clear whether the following definition of the game would be suitable: we assume that the vector $P_n = (P_n^1, \ldots, P_n^m)$ is given and, for any strategy profile s on $[0, T_n]$, we postulate that (4.15) holds and we set, for every $l = 1, \ldots, n-1$,

$$V_l(s_l) := \pi_{\mathbb{H}(\mathcal{E}_l(s_l),X_l,\bar{p}_l(s_{l+1}))}\big(\mathbb{E}_{\mathbb{Q}}\big(V_{l+1}(s_{l+1}) \,|\, \mathcal{F}_{T_l}\big)\big)$$

where, for every $l = 1, \ldots, n-1$,

$$\bar{p}_l(s_{l+1}) := \sum_{i=1}^{m} \mathbb{E}_{\mathbb{Q}}\big(V^i_{l+1}(s_{l+1}) \,|\, \mathcal{F}_{T_l}\big).$$

Unfortunately, this more flexible game is unlikely to possess an optimal equilibrium, in general, and thus its specification should be complemented by further conditions.

Remark 4.2. One may wonder whether the multi-person game introduced in Definition 4.3 can be solved using a reflected BSDE, as was demonstrated by Cvitanić and Karatzas [5] for two-person Dynkin games. This is indeed the case, but it requires a judicious specification of the behaviour of a solution to a multi-dimensional BSDE at the boundary. In [13], it was shown that the BSDE approach can be applied in a discrete-time setting to solve the so-called *multi-player affine games*, which cover the game of Definition 4.3 as a special case. The interested reader is also referred to [29] where continuous-time multi-player stopping games are solved using multi-dimensional BSDEs with oblique reflection.

Acknowledgement. The research of Ivan Guo and Marek Rutkowski was supported under Australian Research Council's Discovery Projects funding scheme (DP120100895).

References

[1] Ayache, E., Forsyth, P. and Vetzal, K.: Valuation of convertible bonds with credit risk. *J. Derivatives*, Fall 2003.

[2] Andersen, L. and Buffum, L.: Calibration and implementation of convertible bond models. *J. Comput. Finance* 7 (2004), 1–34.

[3] Bielecki, T. R., Crépey, S., Jeanblanc, M. and Rutkowski, M.: Arbitrage pricing of defaultable game options with applications to convertible bonds. *Quant. Finance* 8 (2008), 795–810.

[4] Bielecki, T. R. and Rutkowski, M.: Valuation and hedging of contracts with funding costs and collateralization. Working paper, 2013.

[5] Cvitanić, J. and Karatzas, I.: Backward stochastic differential equations with reflection and Dynkin games. *Ann. Probab.* 24 (1996), 2024–2056.

[6] Dolinsky, Y. and Kifer, Y.: Hedging with risk for game options in discrete time. *Stochastics: Int. J. Probab. Stoch. Process.* 79 (2007), 169–195.

[7] Dolinsky, Y., Iron, Y. and Kifer, Y.: Perfect and partial hedging for swing game options in discrete time. *Math. Finance* 21 (2011), 447–474.

[8] Dynkin, E. B.: Game variant of a problem on optimal stopping. *Soviet Math. Dokl.* 10 (1969), 270–274.

[9] Föllmer, H. and Schied, A.: *Stochastic Finance: An Introduction in Discrete Time.* 2nd ed., De Gruyter, Berlin, 2004.

[10] Guo, I.: Competitive multi-player stochastic games with applications to multi-person financial contracts. Doctoral dissertation, University of Sydney, 2013.

[11] Guo, I. and Rutkowski, M.: A zero-sum competitive multi-player game. *Demonstratio Math.* 45 (2012), 415–433.

[12] Guo, I. and Rutkowski, M.: Discrete-time multi-player stopping and quitting games with redistribution of payoffs. In: *Arbitrage, Credit and Informational Risks*, C. Hillairet, M. Jeanblanc and Y. Jiao, eds., World Scientific, Singapore, 2014, pp. 171–206.

[13] Guo, I. and Rutkowski, M.: Stochastic multi-player competitive games in discrete time. Working paper, University of Sydney, 2013.

[14] Hamadène, S. and Hassani, M.: The multi-player nonzero-sum Dynkin game in continuous time. Working paper, Université du Maine, 2011.

[15] Hamadène, S. and Hassani, M.: The multi-player nonzero-sum Dynkin game in discrete time. *Math. Meth. Oper. Res.* 79 (2014), 179–194.

[16] Hamadène, S. and Zhang, J.: The continuous time nonzero-sum Dynkin game problem and application in game options. *SIAM J. Control Optim.* 48 (2010), 3659–3669.

[17] Jacod, J. and Shiryaev, A. N.: Local martingales and the fundamental asset pricing theorems in the discrete-time case. *Finance Stoch.* 2 (1998) 259–273.

[18] Jeanblanc, M., Yor, M. and Chesney, M.: *Mathematical Methods for Financial Markets.* Springer, 2009.

[19] Kallenberg, O.: *Foundations of Modern Probability.* Springer, 1997.

[20] Kallsen, J. and Kühn, C.: Pricing derivatives of American and game type in incomplete markets. *Finance Stoch.* 8 (2004), 261–284.

[21] Kallsen, J. and Kühn, C.: Convertible bonds: financial derivatives of game type. In: *Exotic Option Pricing and Advanced Lévy Models*, Wiley, Chichester New York, 2005, pp. 277–291.

[22] Karatzas, I. and Li, Q.: BSDE approach to non-zero-sum differential games of control and stopping. Working paper, Columbia University.

[23] Karatzas, I. and Shreve, S.: *Methods of Mathematical Finance.* Springer, 1998.

[24] Kifer, Y.: Game options. *Finance Stoch.* 4 (2000), 443–463.

[25] Kifer, Y.: Dynkin games and Israeli options. *ISRN Probability and Statistics* (2013), ID856458, 17 pages.

[26] Kühn, C., Kyprianou, A. E. and van Schaik, K.: Pricing Israeli options: a pathwise approach. *Stochastics: Int. J. Probab. Stoch. Process.* 79 (2006), 117–137.

[27] Kyprianou, A. E.: Some calculations for Israeli options. *Finance Stoch.* 8 (2004), 73–86.

[28] Laraki, R. and Solan, E.: Equilibrium in two-player non-zero-sum Dynkin games in continuous time. *Stochastics: Int. J. Probab. Stoch. Process.* 85 (2013), 997–1014.

[29] Nie, T. and Rutkowski, M.: Multi-player stopping games with redistribution of payoffs and BSDEs with oblique reflection. *Stoch. Process. Appl.* 124 (2014), 2672–2698.

[30] Ohtsubo, Y.: A nonzero-sum extension of Dynkin's stopping problem, *Math. Oper. Res.* 12 (1987), 277–296.

[31] Pallavicini, A., Perini, D. and Brigo, D.: Funding, collateral and hedging: uncovering the mechanism and the subtleties of funding valuation adjustments. Working paper, 2012.

[32] Peskir, G.: Optimal stopping games and Nash equilibrium. *Theory Probab. Appl.* 53 (2008), 623–638.

[33] Snell, J.L.: Applications of martingale system theorems. *Trans. Amer. Math. Soc.* 73 (1952), 293–312.

5 Appendix: Nash and Optimal Equilibria

We summarise here the basic results for Nash and optimal equilibria in a multi-player game. Let $\mathcal{M} = \{1, \dots, m\}$ be the set of players and let $\mathcal{S}$ stand for the class of all strategy profiles $s = (s^1, \dots, s^m)$. For each $s \in \mathcal{S}$, we denote by $V^k(s) = V^k(s^1, \dots, s^m)$ a (possibly random) payoff of the kth player. It is convenient to write $s = (\sigma^k, \sigma^{-k})$ and $\mathcal{M}^{-k} = \mathcal{M} \setminus \{k\}$. An analogous notational convention will be later applied to any proper subset $\mathcal{N}$ of $\mathcal{M}$.

Definition 5.1. A strategy profile $\sigma = (\sigma^1, \dots, \sigma^m) \in \mathcal{S}$ is called a *Nash equilibrium* if the inequality $V^k(\sigma^k, \sigma^{-k}) \geq V^k(s^k, \sigma^{-k})$ holds for each k and all $s^k \in \mathcal{S}^k$. In other words, for each $k \in \mathcal{M}$,

$$V^k(\sigma^k, \sigma^{-k}) = \operatorname*{ess\,sup}_{s^k \in \mathcal{S}^k} V^k(s^k, \sigma^{-k}). \tag{5.1}$$

Often (5.1) is written as $V^k(\sigma^k, \sigma^{-k}) \geq V^k(s^k, \sigma^{-k})$ for all $s^k \in \mathcal{S}^k$.

Let us now recall the notion of an optimal equilibrium, which is stronger than a by far more widely used concept of a Nash equilibrium (although the two concepts coincide in the case of a two-person zero-sum game).

Definition 5.2. A Nash equilibrium $\sigma = (\sigma^1, \dots, \sigma^m) \in \mathcal{S}$ is called an *optimal equilibrium* if the inequality $V^k(\sigma^k, \sigma^{-k}) \leq V^k(\sigma^k, s^{-k})$ holds for each $k \in \mathcal{M}$ and all $s^{-k} \in \mathcal{S}^{-k}$ or, equivalently, for each $k \in \mathcal{M}$,

$$V^k(\sigma^k, \sigma^{-k}) = \operatorname*{ess\,inf}_{s^{-k} \in \mathcal{S}^{-k}} V^k(\sigma^k, s^{-k}). \tag{5.2}$$

Combining with condition (5.1) of a Nash equilibrium, σ satisfies

$$V^k(\sigma^k, \sigma^{-k}) = \operatorname*{ess\,inf}_{s^{-k} \in \mathcal{S}^{-k}} V^k(\sigma^k, s^{-k}) = \operatorname*{ess\,sup}_{s^k \in \mathcal{S}^k} V^k(s^k, \sigma^{-k}), \tag{5.3}$$

or equivalently,

$$V^k(\sigma^k, s^{-k}) \geq V^k(\sigma^k, \sigma^{-k}) \geq V^k(s^k, \sigma^{-k}), \quad \forall\, s^k \in \mathcal{S}^k, s^{-k} \in \mathcal{S}^{-k}, \tag{5.4}$$

for each $k \in \mathcal{M}$.

It is clear that an optimal equilibrium is essentially a saddle point. In addition to the properties of a Nash equilibrium, each player can guarantee his optimal equilibrium payoff without knowing the actions of other players. Furthermore, as shown in Corollary 5.1 below, all optimal equilibria achieve the same value. The concept of an optimal equilibrium is also closely related to the maximin and minimax values of the game. The *maximin value* is the maximum payoff player k can guarantee.

Definition 5.3. The *maximin value* $\underline{V}^k$ equals

$$\underline{V}^k := \operatorname*{ess\,sup}_{s^k \in \mathcal{S}^k} \operatorname*{ess\,inf}_{s^{-k} \in \mathcal{S}^{-k}} V^k(s^k, s^{-k}). \tag{5.5}$$

A *maximin strategy* for player k is any strategy $\widehat{s}^k \in \mathcal{S}^k$ such that $\operatorname{ess\,inf}_{s^{-k} \in \mathcal{S}^{-k}} V^k(\widehat{s}^k, s^{-k}) = \underline{V}^k$, that is, a strategy $\widehat{s}^k$ realises the supremum in (5.5).

The *minimax value* is the lowest payoff that the other players can force upon player k.

Definition 5.4. The *minimax value* $\overline{V}^k$ equals

$$\overline{V}^k := \operatorname*{ess\,inf}_{s^{-k} \in \mathcal{S}^{-k}} \operatorname*{ess\,sup}_{s^k \in \mathcal{S}^k} V^k(s^k, s^{-k}). \tag{5.6}$$

A *minimax strategy profile* for the player set $\mathcal{M}^{-k}$ is any strategy profile $\widehat{s}^{-k} \in \mathcal{S}^{-k}$ such that the equality $\operatorname{ess\,sup}_{s^k \in \mathcal{S}^k} V^k(s^k, \widehat{s}^{-k}) = \overline{V}^k$ holds, that is, $\widehat{s}^{-k}$ realises the infimum in (5.6).

5.1 Value of a Game

In general, the maximin value never higher than the minimax value, since the players from the set $\mathcal{M}^{-k}$ cannot force the payoff of player k to be lower than an amount that can be guaranteed by him. For this reason, the maximin value (minimax value, resp.) is also known as the *lower value* (the *upper value*, resp.). The most fundamental properties of an optimal equilibrium are summarised in the following proposition.

Proposition 5.1. *(i) The inequality $\overline{V}^k \geq \underline{V}^k$ is valid for all k.*
(ii) If σ is a Nash equilibrium, then $V^k(\sigma) \geq \overline{V}^k$ for all k.
(iii) If σ satisfies (5.2), then for all k, $V^k(\sigma) \leq \underline{V}^k$.
(iv) If σ is an optimal equilibrium, then $V^k(\sigma) = \overline{V}^k = \underline{V}^k$ for all k.
(v) If σ is an optimal equilibrium, then for all k the strategy σ^k (the strategy profile σ^{-k}, resp.) is a maximin strategy for player k (a minimax strategy profile the players set $\mathcal{M}^{-k}$, resp.).

Proof. (i) For every $\widehat{s}^k \in \mathcal{S}^k$ and $\widehat{s}^{-k} \in \mathcal{S}^{-k}$, we have that

$$G(\widehat{s}^{-k}) := \sup_{s^k \in \mathcal{S}^k} V^k(s^k, \widehat{s}^{-k}) \geq V^k(\widehat{s}^k, \widehat{s}^{-k}) \geq \inf_{s^{-k} \in \mathcal{S}^{-k}} V^k(\widehat{s}^k, s^{-k}) =: H(\widehat{s}^k)$$

and thus $G(s^{-k}) \geq H(s^k)$ for every s^k and s^{-k}. Consequently,

$$\begin{aligned}\overline{V}^k &= \inf_{s^{-k} \in \mathcal{S}^{-k}} \sup_{s^k \in \mathcal{S}^k} V^k(s^k, s^{-k}) = \inf_{s^{-k} \in \mathcal{S}^{-k}} G(s^{-k}) \\ &\geq \sup_{s^k \in \mathcal{S}^k} H(s^k) = \sup_{s^k \in \mathcal{S}^k} \inf_{s^{-k} \in \mathcal{S}^{-k}} V^k(s^k, s^{-k}) = \underline{V}^k.\end{aligned}$$

(ii) If condition (5.1) holds then

$$V^k(\sigma^k, \sigma^{-k}) = \sup_{s^k \in \mathcal{S}^k} V^k(s^k, \sigma^{-k}) \geq \inf_{s^{-k} \in \mathcal{S}^{-k}} \sup_{s^k \in \mathcal{S}^k} V^k(s^k, s^{-k}) = \overline{V}^k.$$

(iii) If condition (5.2) holds then

$$V^k(\sigma^k, \sigma^{-k}) = \inf_{s^{-k} \in \mathcal{S}^{-k}} V^k(\sigma^k, s^{-k}) \leq \sup_{s^k \in \mathcal{S}^k} \inf_{s^{-k} \in \mathcal{S}^{-k}} V^k(s^k, s^{-k}) = \underline{V}^k.$$

(iv) If both conditions (5.1) and (5.2) hold, then by (ii) and (iii),

$$\underline{V}^k \geq V^k(\sigma) \geq \overline{V}^k.$$

In view of (i), we thus obtain the equality $V^k(\sigma) = \overline{V}^k = \underline{V}^k$.
(v) Combining (iv) with condition (5.4), we have,

$$\begin{aligned}V^k(\sigma^k, s^{-k}) &\geq V^k(\sigma^k, \sigma^{-k}) = \underline{V}^k, \quad \forall\, s^{-k} \in \mathcal{S}^{-k}, \\ V^k(s^k, \sigma^{-k}) &\leq V^k(\sigma^k, \sigma^{-k}) = \overline{V}^k, \quad \forall\, s^k \in \mathcal{S}^k.\end{aligned}$$

Hence σ^k and σ^{-k} are maximin and minimax strategies, respectively. □

Proposition 5.1 motivates the following definition of values of the game.

Definition 5.5. If the equality $\overline{V}^k = \underline{V}^k$ holds, then $V^{*k} := \overline{V}^k = \underline{V}^k$ is the *value for player k*. The *value of the game* V^* is the vector $(V^{*1}, \ldots, V^{*m})$ of values for all players.

Since the equality is not necessarily achieved in part (i) in Proposition 5.1, the existence of the value is not guaranteed. However, by part (iv) in this proposition, the existence of an optimal equilibrium implies the existence of the value of the game, so that we may state the following corollary.

Corollary 5.1. *(i) If the value of the game V^* exists, then it is unique.*
(ii) If there exists an optimal equilibrium σ, then the value exists for every player and

$$V(\sigma) = (V^{*1}, \ldots, V^{*m}) = V^*.$$

Hence every optimal equilibrium σ achieves the same value.

Proof. (i) is implicit in Definition 5.5 and (ii) follows immediately from Proposition 5.1 (iv). Finally (iii) follows immediately from (i) and (ii). □

5.2 Coalition Values

The maximin and minimax values, as well as strategy profiles, can be analogously defined for any proper subset $\mathcal{N}$ of the set $\mathcal{M}$ of players, by simply replacing k with $\mathcal{N}$. For any subset $\mathcal{N} \subset \mathcal{M}$, we denote $V^{\mathcal{N}}(s) := \sum_{i\in\mathcal{N}} V^i(s)$ and we write $s = (s^{\mathcal{N}}, s^{-\mathcal{N}})$ for any strategy profile $s \in \mathcal{S}$.

Let us now focus on a special case of a zero-sum game. It is worth noting that even in zero-sum games, not all Nash equilibria are optimal equilibria.

Definition 5.6. A game is called *zero-sum* if $V^{\mathcal{M}}(s) := \sum_{i\in\mathcal{M}} V^i(s) = 0$ for all $s \in \mathcal{S}$.

In the case of a zero-sum game, further properties an optimal equilibrium can be derived.

Proposition 5.2. *Suppose that the game is zero-sum and denote by σ any strategy profile. Then the following statements are equivalent:*
(i) For each $k \in \mathcal{M}$, $V^k(\sigma^k, \sigma^{-k}) \le V^k(\sigma^k, s^{-k})$ for all $s^{-k} \in \mathcal{S}^{-k}$.
(ii) For any proper subset $\mathcal{N} \subset \mathcal{M}$, $V^{\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \ge V^{\mathcal{N}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}})$ for all $s^{\mathcal{N}} \in \mathcal{S}^{\mathcal{N}}$.
(iii) For any proper subset $\mathcal{N} \subset \mathcal{M}$, $V^{\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \le V^{\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}})$ for all $s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}$.
(iv) The strategy profile σ is an optimal equilibrium.

Proof. [(i) $\Longrightarrow$ (ii)] Intuitively, if each player $k \in \mathcal{N}$ can guarantee his payoff by using σ^k, then collectively, the players of $\mathcal{N}$ can guarantee their total payoff by playing $\sigma^{\mathcal{N}}$. Since for all $s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}$ we have $(\sigma^{\mathcal{N}\setminus\{k\}}, s^{-\mathcal{N}}) \in \mathcal{S}^{-k}$. By (i),

$$V^k(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) = V^k(\sigma^k, \sigma^{\mathcal{N}\setminus\{k\}}, \sigma^{-\mathcal{N}}) \le V^k(\sigma^k, \sigma^{\mathcal{N}\setminus\{k\}}, s^{-\mathcal{N}}) = V^k(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}). \tag{5.7}$$

Summing (5.7) over $k \in \mathcal{N}$, we have

$$V^{\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) = \sum_{k\in\mathcal{N}} V^k(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \le \sum_{k\in\mathcal{N}} V^k(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}) = V^{\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}),$$

as required.

[(ii) $\Longrightarrow$ (i)] This is immediate by setting $\mathcal{N} = \{k\}$.

[(ii) $\Longleftrightarrow$ (iii)] Since the game is zero-sum, $V^{\mathcal{N}}(s) = -V^{-\mathcal{N}}(s)$ for all $s \in \mathcal{S}$. Hence

$$\begin{aligned} & V^{\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \le V^{\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}), && \forall\, s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}, \\ \Longleftrightarrow \quad & -V^{-\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \le -V^{-\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}), && \forall\, s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}, \\ \Longleftrightarrow \quad & V^{-\mathcal{N}}(\sigma^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \ge V^{-\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}), && \forall\, s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}. \end{aligned} \tag{5.8}$$

Relabelling $-\mathcal{N}$ as $\mathcal{N}$ in (5.8) gives the desired result.

[(i) and (iii) $\Longrightarrow$ (iv)] From Definition 5.2, the definition of an optimal equilibrium, it is sufficient to check (5.1)

$$V^k(\sigma^k, \sigma^{-k}) \ge V^k(s^k, \sigma^{-k}), \quad \forall\, s^k \in \mathcal{S}^k$$

and (5.2)

$$V^k(\sigma^k, s^{-k}) \ge V^k(\sigma^k, \sigma^{-k}), \quad \forall\, s^{-k} \in \mathcal{S}^{-k}.$$

This is clear as (i) is identical to (5.1), while (iii) reduces to (5.2) after setting $\mathcal{N} = \{k\}$.

[(iv) $\Longleftrightarrow$ (i)] This is immediate from (5.2) in Definition 5.2. □

Note that the equivalence between Proposition 5.2 (i) and (ii) does not use the fact that the game is zero-sum. Proposition 5.2 (iii) implies (5.1), the Nash equilibrium condition, but the converse does not hold.

So far, the definition of the value referred to the value of the game to each individual player. But suppose some subset of players $\mathcal{N} \subseteq \mathcal{M}$ is playing as a *coalition* with the goal to maximise the *coalition payoff* $V^{\mathcal{N}}(s) := \sum_{i\in\mathcal{N}} V^i(s)$. The definition of the value of the game can be extended as follows.

Definition 5.7. The *value* $V^{*\mathcal{N}}$ *for the coalition* $\mathcal{N}$ of players is given by

$$V^{*\mathcal{N}} := \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}) = \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}),$$

assuming that the second equality holds.

In general, the value $V^{*\mathcal{N}}$ (if well defined) does not necessarily satisfy the *additivity property* $V^{*\mathcal{N}} = \sum_{i\in\mathcal{N}} V^{*i}$. However, as the following propositions shows, if an optimal equilibrium exist in a zero-sum game, then this property is indeed satisfied for any subset $\mathcal{N}$ of $\mathcal{M}$.

Proposition 5.3. *Suppose the game has an optimal equilibrium* $\sigma \in \mathcal{S}$ *with value* $V^* = V(\sigma)$*. If either (i) the game is zero-sum or (ii)* $V^{\mathcal{M}}(\sigma) = \operatorname{ess\,sup}_{s\in\mathcal{S}} V^{\mathcal{M}}(s)$*, then for all subsets* $\mathcal{N} \subseteq \mathcal{M}$*, we have that*

$$V^{\mathcal{N}}(\sigma) = \sum_{i\in\mathcal{N}} V^{*i} = V^{*\mathcal{N}} = \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}) = \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}).$$

In particular, the value of the game for any coalition $\mathcal{N}$ *of players is equal to the sum of values for players from* $\mathcal{N}$*.*

Proof. Since case (ii) includes the zero-sum case (i), it suffices to establish the result for case (ii) only. Since σ is an optimal equilibrium, each player $i \in \mathcal{M}$ can guarantee the payoff $V^i(\sigma)$, in other words,

$$V^i(\sigma) = \inf_{s^{-i}\in\mathcal{S}^{-i}} V^i(\sigma^i, s^{-i}).$$

So the players from $\mathcal{N}$ can play to guarantee $V^{\mathcal{N}}(\sigma)$ by playing $\sigma^{\mathcal{N}}$,

$$V^{\mathcal{N}}(\sigma) = \sum_{i\in\mathcal{N}} \inf_{s^{-i}\in\mathcal{S}^{-i}} V^i(\sigma^i, s^{-i}) \le \sum_{i\in\mathcal{N}} V^i(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}) = V^{\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}), \quad \forall\, s^{-\mathcal{N}} \in \mathcal{S}^{-\mathcal{N}}.$$

Hence

$$V^{\mathcal{N}}(\sigma) \le \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} V^{\mathcal{N}}(\sigma^{\mathcal{N}}, s^{-\mathcal{N}}) \le \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}). \tag{5.9}$$

Applying the same argument to the player set $-\mathcal{N}$, they can also guarantee their payoff $V^{-\mathcal{N}}(\sigma)$,

$$V^{-\mathcal{N}}(\sigma) \le V^{-\mathcal{N}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}}), \quad \forall\, s^{\mathcal{N}} \in \mathcal{S}^{\mathcal{N}}.$$

Now by condition (ii) $V^{\mathcal{M}}(\sigma) = \sup_{s\in\mathcal{S}} V^{\mathcal{M}}(s)$, we must have, for all $s^{\mathcal{N}} \in \mathcal{S}^{\mathcal{N}}$,

$$V^{\mathcal{N}}(\sigma) = V^{\mathcal{M}}(\sigma) - V^{-\mathcal{N}}(\sigma) \ge V^{\mathcal{M}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}}) - V^{-\mathcal{N}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}}) = V^{\mathcal{N}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}}).$$

Hence

$$V^{\mathcal{N}}(\sigma) \ge \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, \sigma^{-\mathcal{N}}) \ge \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}). \tag{5.10}$$

Finally, the required result follows immediately by combining (5.9) and (5.10) with the following fact (see 5.1 (i)),

$$\inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}) \ge \sup_{s^{\mathcal{N}}\in\mathcal{S}^{\mathcal{N}}} \inf_{s^{-\mathcal{N}}\in\mathcal{S}^{-\mathcal{N}}} V^{\mathcal{N}}(s^{\mathcal{N}}, s^{-\mathcal{N}}),$$

completing the proof. □